


**SPECIAL ISSUE: SUPERCOMPUTERS**

IEEE  
**SPECTRUM**

# **TERAFLOPS GALORE**



SEPTEMBER 1992

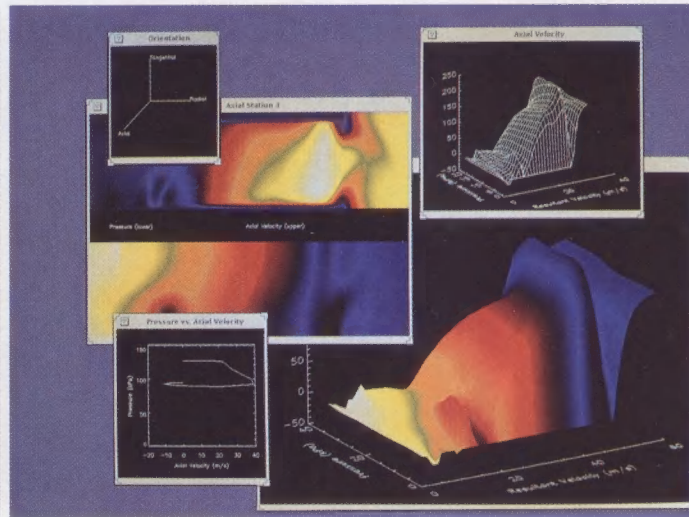



THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.



# RAW.

# WELL DONE.



Attempting to analyze stacks of raw data can be frustrating if they can't be transformed into meaningful information that you can easily understand and use. Trends, relationships, insights, even breakthroughs, can remain buried unless you can unlock your data.

The key is PV-WAVE™, and The Visual Data Analysis Software from Precision Visuals.

PV-WAVE combines graphics with image processing, surface rendering, and animation, allowing you to visually analyze and interpret your data.

## PV-WAVE™ TRANSFORMS DATA INTO DISCOVERY



*"PV-WAVE gives me the power and flexibility of FORTRAN with the ease of use of my spreadsheet and a wide array of graphics capabilities."*

Bill Estell  
Analytical Engineer  
United Technologies Hamilton Standard

With PV-WAVE, you can actually see your data, navigate through it, and display meaningful results. And, depending on your application, choose from point and click or command language versions for your Sun, DEC, HP, IBM, and SGI workstations and multi-user systems running UNIX, ULTRIX, and VMS.

If you're a scientist, engineer, researcher, or analyst who needs to turn raw data into well done results, call us today and ask about our no-risk, 60-day money-back guarantee!

**Call Chris Logan today for your free video demo or evaluation copy of PV-WAVE.**

## 1-800-447-7147



### PV-WAVE

*The Visual Data Analysis Software*

Precision Visuals, Inc. 6230 Lookout Road Boulder, Colorado 80301 USA (303) 530-9000

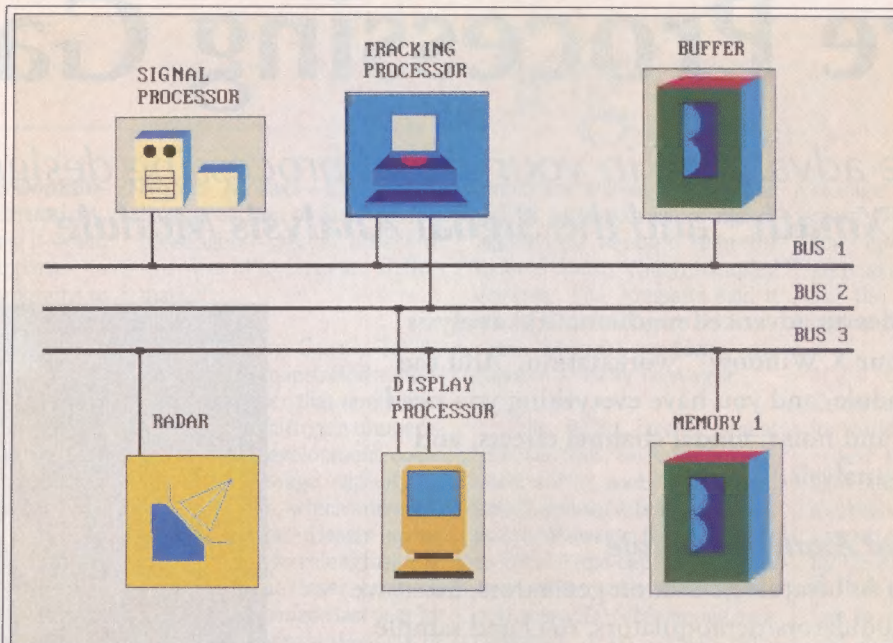
© 1991 Precision Visuals, Inc.

PV-WAVE is a trademark of Precision Visuals, Inc.

Circle No. 31



## NEW FOR NETWORK ANALYSTS AND SYSTEMS ENGINEERS



Realistic simulation of your network or embedded computer system--quick results, no programming

## NETWORK II.5 now predicts performance of computer-communication systems

Free trial and, if you act now, free training

**N**ETWORK II.5 uses simulation to predict your network performance. You simply describe your network and workload.

**Animated simulation follows immediately--no programming delays.**

### Easy-to-understand results

You get an animated picture of your network. System bottlenecks and changing levels of utilization are apparent.

Seeing your network animated increases everyone's understanding of its operation and builds confidence in your results.

Your reports show response times, messages delivered, messages lost, device utilization, and queuing statistics.

### Computers with NETWORK II.5

NETWORK II.5 is available for most PC's, Workstations, and Mainframes.

### Your network simulated

You can analyze embedded or distributed computer systems, or other computer-communication networks. Industry standard protocols such as FDDI and IEEE Standard 802.X are built-in. Others can be modeled.

You can easily study the effect of changing network parameters or even network protocols.

You can simulate some portions of the network at a detailed level and others at a coarser level.

### Free trial information

The free trial contains everything you need to try NETWORK II.5® on your computer. For a limited time we also include free training --no cost, no obligation.

Call Paul Gorman at (619) 457-9681, Fax (619) 457-1184. In Europe, call Joe Lenz, in The Netherlands, on 31 43 670780, Fax 31 43 670200. In the UK, call Nigel McNamara on 0276 671 671, Fax 0276 670 677. In Canada, call Peter Holt on (613) 782-2474, Fax (613) 782-2202.

### Free trial offer

☐ Yes I want to see how NETWORK II.5 quickly answers network performance questions.

**Limited offer--Act now for free training.**

Name

Organization

Address

City

State

Zip

Telephone

Fax

Computer

Op. Syst.

☐ Send details on your University Offer.

**Return to:**

CACI Products Company  
3333 North Torrey Pines Court  
La Jolla, California 92037  
Call Paul Gorman at (619) 457-9681 Mic spec  
Fax (619) 457-1184

**In Canada:**

CACI Products Company  
200-440 Laurier Avenue West  
Ottawa, Ontario K1R 7X6  
Call Peter Holt on (613) 782-2474  
Fax (613) 782-2202

**In Europe:**

CACI Products Division  
MECC Business Center  
G. Martino Laan 85, 8th Floor  
6229GS Maastricht, The Netherlands  
Tel. 31 43 670780, Fax 31 43 670200

**In the UK:**

CACI Products Division  
Coliseum Business Centre  
Watchmoor Park, Riverside Way  
Camberley, Surrey GU15 3YL, UK  
Call Nigel McNamara on 0276 671 671  
Fax 0276 670 677

NETWORK II.5 is a registered trademark and service mark of CACI Products Company.  
©1992 CACI Products Company.



# More Processing Gain !

*Gain the advantage in your signal processing designs with  
**Xmath™** and the **Signal Analysis Module™***

Only Xmath provides an advanced mathematical analysis environment for your X Window™ workstation. Add the Signal Analysis Module, and you have everything you need to generate signals and noise, model channel effects, and perform waveform analysis.

## ***SAM—the Signal Analysis Module***

The SAM gives you full featured function generators, extensive noise generators, modulators/demodulators, A/D and sample and hold models, TFRs, spread spectrum signal generators and much more.

## ***Interactive graphics***

Zoom in for signal details, add text, change colors and viewing angles, or draw in shapes and arrows.

## ***Object-oriented MathScript™***

More than just matrices! Xmath's unique object-oriented language gives you fast execution and frees you to function at a higher level of data abstraction.

## ***Programmable GUI***

Xmath's new programmable GUI layer lets you develop tools which take full advantage of the X Window interface, with no C or Fortran programming.

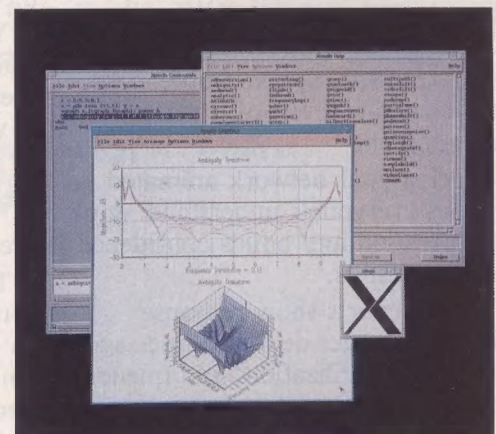
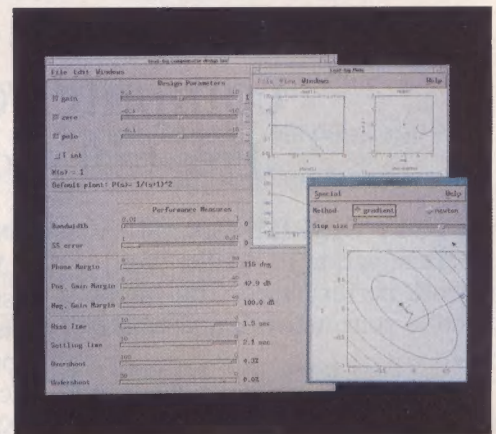

## ***Intuitive User Interface***

Pull-down menus, variable windows, object editors, and the Hypertext Help utility make Xmath easy to learn and use.

***For more gain in your signal processing designs give us a call at 1-800-932-MATH and we'll send you a FREE Xmath/SAM Demo for X Window workstations.***



3260 Jay Street • Santa Clara, California 95054  
Tel: 408.980.1500 • Fax: 408.980.0400  
e-mail: xmath-info@isi.com



© 1992 Integrated Systems, Inc. All rights reserved. Xmath and MathScript are trademarks of Integrated Systems, Inc. Signal Analysis Module is a trademark of ELANIX, Inc. X Window is a trademark of the Massachusetts Institute of Technology.

Object-oriented Analysis For Signal Processing. **Xmath**



## Newslog

**JULY 16.** The **Federal Communications Commission**, Washington, D.C., said it would allow telephone companies to transmit television programs to homes over their phone lines and to buy up to 5 percent of a cable company outside their service areas as well as 5 percent of any video programming company. It also opened up new wireless markets by authorizing more advanced pocket-telephone services.

**JULY 20.** **Hughes Power Control Systems**, Torrance, Calif., said it had developed an inductive charging system for electric cars that eliminates the possibility of shocks even in rainy or damp weather. It can do a routine daily recharge in less than 1 hour, a complete recharge in 2-3 hours. Instead of an electric connector, the system uses a plastic-covered inductive paddle about 125 mm in diameter inserted into a slot on the car. The system determines how much charge the car requires, sets the charging rate, and even tells the local electric company whom to bill.

**JULY 21.** The head of **Russia's Ministry of Atomic Energy** said the republic had discussed with **U.S. Department of Energy** officials the possibility of selling tons of uranium from retired nuclear arms for use in U.S. commercial nuclear power plants. The United States' 110 operating nuclear reactors need at least partial refueling every 12-18 months. The Russians have reportedly also discussed such sales with France and Japan.

**JULY 21.** **British Airways PLC**, London, and **USAir**, Arlington, Va., announced plans to form the world's largest airline partnership, with the British carrier investing US \$750 million to buy up to 44 percent of USAir. The agreement gives USAir the backing to challenge the dominant U.S. airlines, and

British Airways—the world's largest carrier of international passengers—greater access to the world's largest airline market.

**JULY 22.** **Sony Corp.**, Tokyo, said it had demonstrated a semiconductor laser that emits blue light at liquid nitrogen temperatures. The development could triple the storage capacity of compact discs, which currently use red or infrared lasers having much longer wavelengths. Experts said the laser built on breakthroughs made last year by U.S. researchers at the 3M Co. and at Brown and Purdue universities.

**JULY 22.** The **U.S. House of Representatives** announced it had approved a bill to regulate prices for cable television. The measure, similar to one passed by the Senate in January, would allow cities once again to regulate cable rates, and would force big cable companies with affiliates that produce programming to make their offerings available to competing distributors.

**JULY 24.** The **National Aeronautics and Space Administration (NASA)** said it had launched a rocket with a Japanese scientific satellite that is to swing by the moon and be flung by lunar gravity about 1.5 million kilometers into space. The goal: better understanding and prediction of the effects of the solar wind (charged particles flowing from the sun) on the tail of the earth's comet-shaped magnetosphere. The US \$160 million, four-year mission is a joint project between NASA and the **Japanese Institute for Space and Astronautical Science**.

**JULY 25.** **Mitsubishi Electric Corp.**, Osaka, Japan, said it has developed a sensor that can read color images with an accuracy of 16 dots per millimeter, reading a page in 10 seconds, not the 2

minutes of current sensors. The use of phototransistors makes production simpler than for models using charge-coupled devices. The company said it would use the device in scanners, facsimile machines, and copiers by early next year.

**JULY 25.** **KDD**, Japan's largest international telecommunications carrier, said it had agreed with **Korean Telecom**, Seoul, and the **Russian Government** to build large-capacity marine optical-fiber cable links capable of sending 560 Mb/s among their countries. Two Danish telecommunications firms, now using cables under the Sea of Japan, will also join the US \$160 million project. Once completed, the project will provide the equivalent of 7560 telephone lines between Japan and Russia.

**JULY 29.** **AT&T**, **General Electric**, **Honeywell**, and **IBM** announced that they were forming a consortium to work on optoelectronics technology that will enhance future computers. The group hopes to adapt current optical-fiber technology for data transmission at rates of at least 500 Mb/s across 10 meters. It further plans to combine 32 of the lines in parallel to form 16-Gb/s links and to operate four of the links in parallel to transfer data at 64 Gb/s.

**AUG 5.** Researchers at **AT&T Bell Laboratories**, Murray Hill, N.J., said they had demonstrated an optical data storage technique that holds nearly 100 times more data than a compact disc and 300 times more than magnetic storage systems. The technique is a spinoff from work they did earlier in a new technology called near-field scanning optical microscopy, which has yielded advances in viewing small objects like living cells.

**AUG 6.** **Sanyo Electric Co.**, Osaka, Japan, said it had developed an orange-light semiconductor laser that could boost the

storage capacity of compact discs. The company said the laser operates at room temperature at a wavelength of 615 nm vs. the 780 nm of red-laser optical-disc equipment.

**AUG 6.** **GTE Corp.**, Stamford, Conn., said it had agreed to sell its worldwide lighting business—best known for its **Sylvania** light bulbs—to **Osram GmbH**, a subsidiary of **Siemens AG**, and a group of international investors for US \$1.1 billion. Analysts said the sale underscored GTE's desire to concentrate on its local and cellular-telephone businesses.

**AUG 7.** **Toshiba Corp.**, Tokyo, and **L.M. Ericsson**, Stockholm, Sweden, said they would form a joint venture to supply digital mobile telecommunications equipment to Japan's **Digital Phone**, a group of regional cellular mobile telephone service companies that will start operation in 1994.

**AUG 9.** **NASA** said the **space shuttle Atlantis** landed in Cape Canaveral, Fla., after releasing into orbit a science satellite for the **European Space Agency**. However, another experiment, one to generate electricity with an **Italian Space Agency** spacecraft tethered 19 km away, was called off by NASA when the crew failed to reel it out more than 0.25 km. The experiment cost US \$379 million and took two decades to plan.

## Preview:

**SEP 24-25.** **Spectrum 20/20 '92**, sponsored by the Radio Advisory Board of Canada and Communications Canada, is to be held in Toronto. Titled "Transitions," the symposium is to assess the impact of WARC '92 and exchange views on issues that will affect future radio spectrum usage. For information, call 613-224-1741.

COORDINATOR: Sally Cahur



# IEEE SPECTRUM

## SPECIAL ISSUE

### SUPERCOMPUTING

#### 26 Introduction

By GLENN ZORPETTE

#### 28 Reinventing the machine

A remarkable diversity of massively parallel machines is ushering in a new era in high-performance computing, *Glenn Zorpette 28*. Huge efforts are going into software for the new generation of supercomputers, *Richard Comerford 34*. A computer paradigm is needed if massively parallel machines are to live up to their billing, *George Cybenko and David J. Kuck 39*.

#### 42 National efforts

The speed champion in uniprocessors, Japan is coming on strong in multiprocessing, *David K. Kahaner and Ulrich Wattenberg 42*. Giant memories and fast networks are cornerstones of the emerging U.S. computing initiative, *Charles E. Catlett 48*.

#### 56 Emerging applications

The challenges awaiting tomorrow's supercomputers run the gamut from the grand to the mundane, *Morris Grossman 56*. Researchers are counting on visualization to help them get the most out of supercomputers, *Matthew Arrott and Sara Latta 61*.

#### 66 Future technologies

In their quest for ever-higher processing rates, supercomputer designers are developing remarkable semiconductor technologies, *Harold Dozier 66* and *Jan Wikstrom and Tony Vacca 68*. Packaging the new semiconductors will also demand cutting-edge technologies, *George F. Watson 69*. Physical limits will force 21st century supercomputer designers to make radical departures, *Trudy E. Bell 72*.


#### 76 To probe further

## SPECTRAL LINES

#### 25 But is it ethical?

By DONALD CHRISTIANSEN



Do engineers behave ethically in depicting a competitor's product as inferior or unpatentable?




Supercomputer image color-codes drag on keel of an America's Cup yacht

Supercomputer visualization of oil reservoir


Parsytec GC supercomputer has up to 16 384 processors







Geometries of organic molecules



Gray Research Inc.

Cognition Inc.

- 3 Newslog
- 6 Reflections
- 8 Calendar
- 12 Books
- 18 The engineer at large
- 22 Forum
- 80 Program notes
- 84 EE's tools & toys
- 94 Scanning THE INSTITUTE
- 94 Coming in *Spectrum*

**Cover:** The 1991 discovery of planets orbiting a pulsar prompted a video animation, produced for television news programs. An enhanced version, from which this scene was taken, shows how the planets cause irregularities in the pulsar's motion and in the timing of radio pulses from the pulsar received on earth. The visualization was done by Wayne Lytle of the Cornell Theory Center, using IBM Corp.'s power visualization system. *Spectrum's* Special Report on Supercomputers begins on p. 26.

IEEE SPECTRUM (ISSN 0018-9235) is published monthly by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. © 1992 by The Institute of Electrical and Electronics Engineers, Inc., 345 East 47th St., New York, N.Y. 10017, U.S.A. Cable address: ITRIPLEE. Telex 236-411. Fax: 212-705-7453. E-mail: ieee@spectrum.

**ANNUAL SUBSCRIPTIONS.** IEEE members: \$11.00 included in dues. Nonmembers: \$29.95. Libraries/institutions: \$139.

**SINGLE COPIES.** Members: \$8 first copy, \$16 per additional copy. Nonmembers: \$16.

**MICROFICHE SUBSCRIPTIONS.** Members: \$16. Nonmembers and libraries: \$139.

**POSTMASTER:** Please send address changes to *IEEE Spectrum*, c/o Coding Department, IEEE Service Center, 445 Hoes Lane, Box 1331, Piscataway, N.J. 08855. Second Class postage paid at New York, N.Y., and additional mailing offices. Canadian GST #125634188.

Printed at 8649 Hacks Cross Rd., Olive Branch, Miss. 38654.

*IEEE Spectrum* is a member of the Audit Bureau of Circulations, the Magazine Publishers of America, and the Society of National Association Publications.

## Staff

**EDITOR AND PUBLISHER:** Donald Christiansen

**ISSUE EDITOR:** Glenn Zorpette

**MANAGING EDITOR:** Alfred Rosenblatt

**SENIOR TECHNICAL EDITOR:** Gadi Kaplan

**SENIOR EDITORS:** Trudy E. Bell, Richard Comerford, Tekla S. Perry, Michael J. Riezenman, George F. Watson

**SENIOR ASSOCIATE EDITORS:** John A. Adam, Glenn Zorpette

**HEADQUARTERS:** New York City 212-705-7555

**BUREAUS:** WASHINGTON, D.C., John A. Adam, 202-544-3790; SAN FRANCISCO, Tekla S. Perry, 415-282-3608

**CORRESPONDENTS:** Fred Guterl, Roger Milne (London); Bradford Smith (Paris); John Blau (Düsseldorf); Robert Ingersoll (Bonn); John Mason (Barcelona); Stuart M. Dambrot, Roger Schreffler (Tokyo); Kim Nak-Hieon (Seoul); Chris Brown (Taipei); Peter Gwynne (Hong Kong); Tony Healy (Sydney, Australia); Christopher Trump (Toronto); Axel de Tristan (Rio de Janeiro); Kevin L. Self (Houston)

**CHIEF COPY EDITOR:** Margaret Eastman

**COPY EDITOR:** Sally Cahur

**EDITORIAL RESEARCHER:** Alan Gardner

**CONTRIBUTING EDITORS:** Karl Esch, Ronald K. Jorgen, Michael F. Wolff

**EDITORIAL SUPPORT SERVICES:**

Rita Holland (Manager)

**EDITORIAL ASSISTANTS:** Ramona Foster, Desiree Noel

**DESIGN CONSULTANT:** Gus Sauter

**OPERATIONS DIRECTOR:** Fran Zappulla

**BUSINESS MANAGER:** Robert T. Ross

**PRODUCTION AND QUALITY CONTROL:** Carol L. White (Director)

**EDITORIAL PRODUCTION:**

Marcia Meyers (Manager)

Peter Ruffett (Typographer)

Morris Khan (Technical Graphic Artist)

**ASSOCIATE PUBLISHER:** William R. Saunders

**ADMINISTRATIVE ASSISTANT:** Carmen Cruz

**MAIL LIST SALES:** Shelly Newman (Manager), Lizette Graciani

**ADVERTISING PRODUCTION:**

Theresa Fitzpatrick (Manager), Francesca Silvestri

**MARKETING DIRECTOR:** Arthur C. Nigro

**PROMOTION MANAGER:** Robert D. Moran

**RESEARCH MANAGER:** Hendrik Prins (Manager), Carl Leibman (Associate)

**MARKETING SERVICES:** Eric Sonntag (Administrator)

**ADMINISTRATIVE ASSISTANT TO THE EDITOR**

**AND PUBLISHER:** Nancy T. Hantman

## Advisory Board

**CHAIRMAN:** G.P. Rodrigue

Charles K. Alexander, B. Leonard Carlson, Donald Fleckenstein, Robert W. Lucky, Irene C. Peden, Cary R. Spitzer, Jerome J. Suran, William R. Tackaberry

## Editorial Board

**CHAIRMAN:** Donald Christiansen

Robert A. Bell, Dennis Bodson, Kjell Carlsen, W. Bernard Carlson, James E. Carnes, Pallab K. Chatterjee, Jacques J. Clade, Robert S. Cooper, Robert P. Davidson, Murray Eden, Alex Hills, Robert R. Johnson, Ted G. Lewis, Michael S.P. Lucas, Tsugio Makimoto, Edith W. Martin, Bruce C. Mather, M. Granger Morgan, David A. Patterson, Alfred R. Potvin, W. David Pricer, Betty Prince, Rob A. Rutenbar, Bruce D. Shriver, Stephen B. Weinstein



# Reflections

## Spacewars

**N**ot long ago I was on a visiting committee enlisted to advise an organization on its most critical problems. "What do you think your most critical problem is?" we asked. "Space," they said. "There isn't enough of it." Of course, they did not have to add the latter clarification; I've never heard anyone complain about having too much space.

One of my fellow committee members, speaking from his infinite wisdom, gave his advice. "Set up a space committee, and survey the available space," he said crisply.

I shook my head slightly and said nothing. Where had I heard this before? Essentially everywhere. Everybody has a space problem, and everybody does surveys. They never work.

Though we do not even admit it to ourselves, we all know in our hearts what happens to space surveys. To ascertain the availability of surplus space, the committee pays a surprise visit to a laboratory—a surprise, that is, to everyone but the people being visited.

There is something about space committees that makes them hard to keep secret. Armed with official clipboards, they approach with all the subtlety of a herd of elephants. The warning jungle noises precede them. Shrill monkey calls are heard, and the cries of hyenas rend the air. Birds twitter, and vultures circle. None of this is missed by the sentries posted by the laboratory to be visited.

The committee chairperson checks the address on the laboratory door against his clipboard. After jotting a small notation on his pad, he knocks tentatively on the door. Sounds of scurrying are heard inside, but there is no other response. He tries pushing the door, but it only opens a few centimeters before it meets some firm resistance. The committee hears a muffled voice from inside. "Just a minute," says the voice, clearly under some physical strain. "We're trying to move this cabinet to let you in."

After a discordant medley of sliding and groaning noises, the door is pushed open enough to admit most of the committee. It is quickly discovered that not all the visitors fit. Those who squeeze in see strange-looking electrical apparatus filling every cranny of the laboratory. The room is littered with scopes, testgear, and racks of circuit boards with wires dangling in every direction. Large pipes are suspended from the

ceiling, and hundreds of cables snake their way into crevices in the raised floor.

The chairperson looks again at his clipboard, on which this lab had been prematurely identified as "empty." He turns to the disheveled person who had opened the door for them. "Are you the only occupant of this facility?" he asks.

The occupant looks puzzled. He turns to look around the room. "Joe," he calls out in a loud voice. "Where are you?"

The committee, too, looks around in puzzlement. But the disheveled occupant continues to call out. "Michael? Mary? Carlos?"

The distinctive bang of a human head against steel is heard, and a face appears from underneath a desk that had itself been previously unnoticed because of the piles of books and catalogs blanketing its surface. "Can't I have any peace and quiet to fill out these equipment order forms?" complains the face peevishly.

The face disappears back beneath the desk. Small, stirring movements throughout the lab create the queasy illusion of a ship at sea as an arm reaches out in one place, a foot moves in another, a face passes between pipes.

"Never mind," stutters the horrified space committee chairperson. "We'll go elsewhere. Sorry to have disturbed you."

Retreating down the corridor, the chairman pens a notation against the lab number on his clipboard: "Fully occupied." Indeed, this same notation is recorded against every lab that has thus far been visited.

Meanwhile there is much celebration back in that "fully occupied" lab. The disheveled

greeter combs his hair and adjusts his tie amid the laughter and handslaps of congratulations. The joy is interrupted by the entrance of the responsible supervisor. "Good job, everybody," he says. "But there's work to be done. This junk has to be moved down to... let's see, room 536," he says, consulting a folded list. "Let's get cracking."

As the space committee takes time for morning coffee and recapitulation of their tally sheets, unbeknownst to them at least part of the stuff they've just seen is rolled on dollies toward their next destination. When casually rearranged, this stuff probably won't be recognized by the committee.

Meanwhile, let us ponder the question: why are we so possessive of even unused labs? After all, in these days of virtual environments whatever you could accomplish in that lab could probably be done more easily through simulation on your desktop computer. But that's hardly the point; this has nothing to do with utility. Everyone knows what is at stake here—*space is power*. The more you have, the more important you are. Clearly, your strategy should be to acquire as much space as possible.

The cardinal rule in space strategy is *never throw away old equipment and experiments*. Keeping your old junk enables you to start on your strategy of acquisition. As you build up more racks of equipment, you will need lab assistants to maintain your growing lab. They, in turn, will require desks and office space, as well as more bench space for their individual sets of test equipment.

Your computing needs soon will exceed the capabilities of the facilities you have been sharing. Thus you will need your own computer center, which will also require the hiring of a systems administrator. The growing power needs will necessitate rewiring of your lab. (In general, the more custom modifications, the better.) By this time the threshold below which personal intervention was needed will have been exceeded; space demands will escalate on their own.


As you acquire more space, it will come to the attention of management. You might naively think that this would count against you. On the contrary, it will be taken very much in your favor. You will be labeled as someone who *needs* space, for the self-evident reason that you have become a star.

Only one thing stands in the way of your stardom—the dreaded space committee. So post your sentries and practice your drills. They will be here any minute. I hear the monkeys shrieking.

Robert W. Lucky


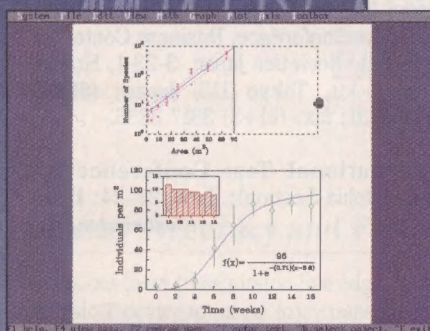




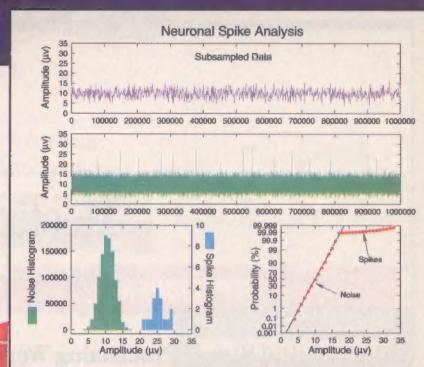


Curve fit in two, three or more dimensions! Create 3-D mesh and scatter plots. Rotate axis and tick labels to match.


With the added feel of a desktop publisher, placing and re-sizing multiple graphs is a snap!



Create, overlay and inset both two- and three-dimensional graphs. Don't forget error bars, axis breaks, and time series tick labels!



Graph huge datasets! This example contains over a million data points. Note data point sampling in top graph.



Slice and rotate your 3-D graphs at will! Section data at any point and take advantage of TRUE 3-D clipping.

# Enter a New Dimension with SigmaPlot® 5.0

**NEW**

SigmaPlot for DOS, the choice of over 30,000 researchers world-wide for creating publication quality graphs, enters a new dimension! SigmaPlot

now includes everything you need to create outstanding 2D and 3D graphs for lectures, poster sessions and journal publication!

## Important Scientific Features

Unlike business graphics packages, SigmaPlot was designed specifically for scientists. With automatic error bars, huge dataset handling, axis breaks, regression lines, confidence intervals, non-linear curve fitting and much more. With superb output and all these specialized features SigmaPlot graphs have been published in hundreds of technical journals!

## Many New Enhancements

With SigmaPlot 5.0 3D plotting is just the beginning!

**Version 4.1 Available for the Mac**

Now you can merge graph files in one easy step; use position and alignment tools for right and left justifying or centering; import SigmaPlot for

Mac® graphs; and refer to improved and expanded documentation. But best of all, creating graphs is easier and faster than ever before.

## Advanced Data Analysis

SigmaPlot's data analysis is among the most advanced found in any graphics package. Input your data directly from the keyboard, spreadsheet and ASCII files into a huge data worksheet. Use SigmaPlot's extensive range of mathematical functions to transform or generate data. Fit nearly any equation to your data with SigmaPlot's powerful non-linear curve fitter.

## Realm of Desktop Publishing

With interactive WYSIWYG page display similar to that found in desktop publishing packages, creating and editing graphs is a snap! Put multiple graphs on a page, place text labels anywhere on the page and move or resize graphs by simply pointing and clicking. SigmaPlot gives you the flexibility and control necessary to customize your graphs down to the last detail! Choose from a large selection of plot types and axis scales, and set axis

ranges, increments, tick marks and grids to your specifications. Create graphs just the way you want them and output to almost any printer, plotter or graphics file.

## Call Today

You have nothing to lose — SigmaPlot comes with an unconditional money-back guarantee and free, unlimited technical support. Call us today and see for yourself why SigmaPlot is the #1 plotting program for scientists and engineers!



**Jandel**  
SCIENTIFIC

2591 Kerner Blvd. San Rafael CA 94901  
PH: 415-453-6700 • FAX: 415-453-7769  
In Europe: Jandel Scientific  
Schimmelbuschstraße 25 • 4006  
Erkrath, Germany  
PH: 02104/36098 • FAX: 02104/33110

**Call for Free Brochure: 800-874-1888**

Circle No. 6 for Mac version.

Circle No. 7 for PC version.



# Calendar

## Meetings, Conferences and Conventions

### SEPTEMBER

**Fifth Digital Signal Processing Workshop (SP);** Sept. 13-16; Starved Rock Lodge, Starved Rock State Park, Ill.; Mark J.T. Smith, School of Electrical Engineering, Georgia Institute of Technology, Atlanta, Ga. 30322-0250; 404-894-6291.

**International Conference on Control and Applications (CS, Dayton Section);** Sept. 13-16; Stouffer Center Plaza Hotel, Dayton, Ohio; Daniel W. Repperger, Armstrong Laboratory, Wright Patterson AFB, Dayton, Ohio 45433-6573; 513-255-5742; fax, 513-255-9687.

**Electrical Overstress/Electrostatic Discharge Symposium (CHMT);** Sept. 15-18; Loew's Anatole Hotel, Dallas; EOS/ESD Association, 200 Liberty Plaza, Rome, N.Y. 13440; 315-339-6937; fax, 315-339-6793.

**Conference on Wireless LAN Implementation (C);** Sept. 17-18; Dayton Convention Center, Dayton, Ohio; IEEE Computer Society, Conference Department, 1730 Massachusetts Ave., N.W., Washington, D.C. 20036-1903; 202-371-1013; fax, 202-728-0884.

**42nd Annual IEEE Broadcast Symposium (BT);** Sept. 17-18; Hotel Washington, Washington, D.C.; Edmund A. Williams, Public Broadcasting Service, 1320 Braddock Place, Alexandria, Va. 22314; 703-739-5172.

**Virtual Reality Annual International Symposium (NN);** Sept. 18-23; Sheraton Hotel, Seattle, Wash.; Thomas Caudell, Boeing Computer Services, Boeing Building 33-07, MS 7L-22, 2760 160 Ave. S.E., Bellevue, Wash. 98008; 206-865-3763.

**Autotestcon '92 (AES, IM, Dayton Section);** Sept. 21-24; Dayton Convention Center, Dayton, Ohio; Kenneth Wilkinson,

Ateam Corp., 7920 Chambersburg Rd., Dayton, Ohio 45424; 513-237-7971.

**Application Specific Integrated Circuits Conference and Exhibit (C, Rochester Section);** Sept. 21-25; Rochester Riverside Convention Center, Rochester, N.Y.; Lynne M. Engelbrecht, ASIC Seminar Coordinator, 170 Mount Read Blvd., Rochester, N.Y. 14611; 716-328-2310; fax, 716-436-9370.

**13th International Semiconductor Laser Conference (LEO);** Sept. 21-25; Takamatsu Kokusai Hotel, Takamatsu, Japan; 13th IEEE International Semiconductor Laser Conference, Business Center for Academic Societies Japan, 3-23-1, Hongo, Bunkyo-ku, Tokyo 113, Japan; (81+3) 3817 5831; fax, (81+3) 3817 5836.

**International Test Conference (C, Philadelphia Section);** Sept. 22-24; Balti-

(Continued on p. 14)

## Not Another Electromagnetics Analysis Software ...

No ... COSMOS/M is not like other analysis software. With COSMOS/M, you can conduct complex 2- and 3D electromagnetic analyses on your PC or workstation at a fraction of the cost of other packages.

Structural Research, a leading developer of desktop finite

element analysis software, makes it easy for you to prepare, analyze and display the results - you may even link the results to linear and nonlinear static and thermal modules.

COSMOS/M's powerful pre- and postprocessor makes modeling and displaying results virtually effortless. COSMOS/M is ideal for quickly evaluating conceptual designs as it

allows you to save time and eliminate expensive prototyping and testing.



What does this mean? It means that if you are designing actuators or electric motors - or anything in between - you will find that COSMOS/M

makes it faster, easier and more affordable.

But you don't have to take our word for it. Call us for a free limited version and analyze one of your own problems to see how easy and affordable analysis can be.




### COSMOS/M Features

- Magnetostatic analysis
- Transient eddy current analysis
- AC eddy current analysis
- Electrostatic analysis
- Current flow
- Nonlinear analyses
- Electro-thermal coupling
- Magneto-structural coupling
- Menu driven pre- and postprocessor
- 2- and 3D automatic meshing
- Electric and magnetic field distributions
- Forces and torques
- Inductances
- Capacitance matrices

**Run your own design problems ...  
Call now for a free limited  
version of COSMOS/M  
310-452-2158**

1661 Lincoln Blvd., Suite 200  
Santa Monica,  
CA 90404 USA  
tel. 310-452-2158  
fax. 310-399-6421





## We brought you PSpice — the world's most popular mixed analog and digital circuit simulator.

Now we bring you the **Design Center** — the universal analog and digital circuit design environment! Offering integrated schematic capture, mixed analog/digital simulation, and graphical waveform analysis, the **Design Center** will take you and your circuit design from conception to realization with ease and efficiency.

For more information about the **Design Center** and its options, complete and mail this postage-paid card today! Please check the box for the platform you are using.

☐ IBM-PC DOS

☐ Macintosh

☐ SPARCstation

Name \_\_\_\_\_

Company \_\_\_\_\_

Address \_\_\_\_\_

City, State, Zip \_\_\_\_\_


Phone \_\_\_\_\_



### MicroSim Corporation

THE MAKERS OF PSpICE

SPEC 9/92



#### Comparison of the ideal and lossy transmission line models

### Greater Accuracy and Faster Simulation

That's our lossy transmission line edge. We've extended PSpice's intrinsic T device to model lossy transmission lines using the distributed model rather than the traditional lumped approach. What you get is ■ smoother approximation of the line's behavior in significantly less time. Here's why.

### The Distributed Model Advantage

Our lossy line behavior is modeled in the frequency domain instead of with the commonly used ladder structures associated with lumped models. Using this technique, lossy lines can be easily specified in terms of their electrical length and the resistance, inductance, capacitance, and conductance distributed along the entire length of the line. Frequency dependent phenomena, such as skin effect and dielectric loss, is straightforward to simulate.

### Elimination of Spurious Oscillations

Our lossy lines are modeled ■ continuous lines. Thus, the simulation artifacts (oscillations) observed in lines modeled as a finite set of lumped segments, vanish with the distributed model.

### Optimized Simulation Performance

For any given circuit, the continuous lines produced by a distributed model avoid the performance overhead created by the multiple line segments required for accurate results using the lumped model approach. Hence, simulation time can be several times faster with the distributed model.

### Model Loss, Dispersion, and Skin Effect with Accuracy and Efficiency

That sums up the PSpice lossy line edge. Discover for yourself why PSpice continues to hold the industry edge on mixed analog and digital circuit simulation.

For more information on PSpice and the **Design Center**, call MicroSim Corporation toll free at (800) 245-3022 or FAX at (714) 455-0554.



### MicroSim Corporation

*The Standard for Circuit Design*

20 Fairbanks • Irvine, CA 92718

THE MAKERS OF PSpICE

PSpice is a registered trademark of MicroSim Corporation

Circle No. 40 for IBM PC.

Circle No. 41 for SPARCstation.

Circle No. 42 for Macintosh.



NO POSTAGE  
NECESSARY  
IF MAILED  
IN THE  
UNITED STATES

SEPTEMBER

**Fifth Digital Signal Processing Workshop (SP);** Sept. 13-14; Starved Rock State Park, J.T. Smith, School of Electrical Engineering, Georgia Institute of Technology, 30322-0250; 404-894-6200.

**International Conference on Electromagnetic Interference and Applications (CS);** Sept. 13-16; Stouffer Center, Dayton, Ohio; Daniel W. Stouffer, Wright-Patterson Air Force Base, Dayton, Ohio 45433-6570; fax, 513-255-9687.

**Electrical Overstress and Discharge Symposium (ESD);** Sept. 15-18; Loew's Anatole Hotel, Dallas; EOS/ESD Association, 200 Liberty Plaza, Rome, N.Y. 13440; 315-339-6937; fax, 315-339-6793.

**BUSINESS REPLY MAIL**

FIRST CLASS MAIL PERMIT NO. 9677 IRVINE, CA

POSTAGE WILL BE PAID BY ADDRESSEE

**MicroSim Corporation**

20 FAIRBANKS  
IRVINE CA 92718-9905



Bellevue, Wash. 98008; 206-865-3763.

**Autotestcon '92 (AES, IM, Dayton Section);** Sept. 21-24; Dayton Convention Center, Dayton, Ohio; Kenneth Wilkinson,

3817 5831; fax, (81+3) 3817 5836.

**International Test Conference (C, Philadelphia Section);** Sept. 22-24; Baltimore, Md. (Continued on p. 14)

## Not Another Electromagnetics Analysis Software ...


No ... COSMOS/M is not like other analysis software. With COSMOS/M, you can conduct complex 2- and 3D electromagnetic analyses on your PC or workstation at a fraction of the cost of other packages.

Structural Research, a leading developer of desktop finite element analysis

software, makes it easy for you to prepare, analyze and display the results - you may even link

the results to linear and nonlinear static and thermal modules. COSMOS/M's powerful pre- and postprocessor makes modeling and displaying results virtually effortless. COSMOS/M is ideal for quickly evaluating conceptual designs as it

allows you to save time and eliminate expensive prototyping and testing.



What does this mean? It means that if you are designing actuators or electric motors - or anything in between - you will find that COSMOS/M

makes it faster, easier and more affordable.


But you don't have to take our word for it. Call us for a free limited version and analyze one of your own problems to see how easy and affordable analysis can be.

### COSMOS/M Features

- Magnetostatic analysis
- Transient eddy current analysis
- AC eddy current analysis
- Electrostatic analysis
- Current flow
- Nonlinear analyses
- Electro-thermal coupling
- Magneto-structural coupling
- Menu driven pre- and postprocessor
- 2- and 3D automatic meshing
- Electric and magnetic field distributions
- Forces and torques
- Inductances
- Capacitance matrices

Run your own design problems...  
Call now for a free limited  
version of COSMOS/M  
310-452-2158


1661 Lincoln Blvd., Suite 200  
Santa Monica,  
CA 90404 USA  
tel. 310-452-2158  
fax. 310-399-6421



Circle No. 21



# Our Lossy Transmission Lines Have the Edge!



Comparison of the ideal and lossy transmission line models

## Greater Accuracy and Faster Simulation

That's our lossy transmission line edge. We've extended PSpice's intrinsic T device to model lossy transmission lines using the distributed model rather than the traditional lumped approach. What you get is a smoother approximation of the line's behavior in significantly less time. Here's why.

## The Distributed Model Advantage

Our lossy line behavior is modeled in the frequency domain instead of with the commonly used ladder structures associated with lumped models. Using this technique, lossy lines can be easily specified in terms of their electrical length and the resistance, inductance, capacitance, and conductance distributed along the entire length of the line. Frequency dependent phenomena, such as skin effect and dielectric loss, is straightforward to simulate.

## Elimination of Spurious Oscillations

Our lossy lines are modeled as continuous lines. Thus, the simulation artifacts (oscillations) observed in lines modeled as a finite set of lumped segments, vanish with the distributed model.


## Optimized Simulation Performance

For any given circuit, the continuous lines produced by a distributed model avoid the performance overhead created by the multiple line segments required for accurate results using the lumped model approach. Hence, simulation time can be several times faster with the distributed model.

## Model Loss, Dispersion, and Skin Effect with Accuracy and Efficiency

That sums up the PSpice lossy line edge. Discover for yourself why PSpice continues to hold the industry edge on mixed analog and digital circuit simulation.

For more information on PSpice and the **Design Center**, call MicroSim Corporation toll free at (800) 245-3022 or FAX at (714) 455-0554.



**MicroSim Corporation**

*The Standard for Circuit Design*

20 Fairbanks • Irvine, CA 92718

**THE MAKERS OF PSpice**

PSpice is a registered trademark of MicroSim Corporation

Circle No. 40 for IBM PC.

Circle No. 41 for SPARCstation.

Circle No. 42 for Macintosh.



# Books

## From a trade school to an institution

Allen E. Puckett

### Millikan's School: A History of the California Institute of Technology.

Goodstein, Judith R.,  
W.W. Norton, New  
York, 1991, 288 pp.,  
\$25.00.



From its humble birth 101 years ago as a sort of nondescript trade school in the small community of Pasadena, Calif., the California Institute of Technology (Caltech) has blossomed into one of the world's leading schools of science and engineering. Its faculty, past and present, includes some 20 Nobel Prize winners, and its graduates have gone on to key roles in Government, industry, and academia. At the outset, no one

could have possibly imagined such an auspicious future.

*Millikan's School* is the story of how three brilliant and strong-willed scientists—George Ellery Hale, Alfred Noyes, and Robert Millikan—brought about this remarkable transformation almost by themselves. It also describes some of the crucial debates of the 1920s and '30s over the role of universities, Government, and the private sector in the conduct of scientific research. These issues—where research could best be done, and how it should or could be supported—are still very much alive today, and the book's coverage of these and other subjects should appeal to an audience beyond those who have had any association with Caltech.

Originally, the school was named Throop University after its founder, the midwestern lumberman Amos Throop. By great good fortune, Hale came to Pasadena in 1903 ■ ■ young astronomer charged with building a solar observatory on nearby Mount Wilson. He became interested in the fledgling school and joined its board of trustees—an event that would prove decisive in shaping

the institute's future. Hale, a talented scientist, organizer, and fund-raiser, was above all a firm believer in the importance of attracting the very best scientific talent in the country to the new school.

His first brilliant success was in persuading the great chemist, Arthur Noyes, to leave the Massachusetts Institute of Technology (MIT) for Pasadena in 1916. An important factor in Noyes' decision was Hale's conviction that engineering education must be firmly rooted in basic sciences, which Noyes evidently felt was not the case at MIT. Together these two men defined the educational philosophy that distinguished Throop, and subsequently Caltech.

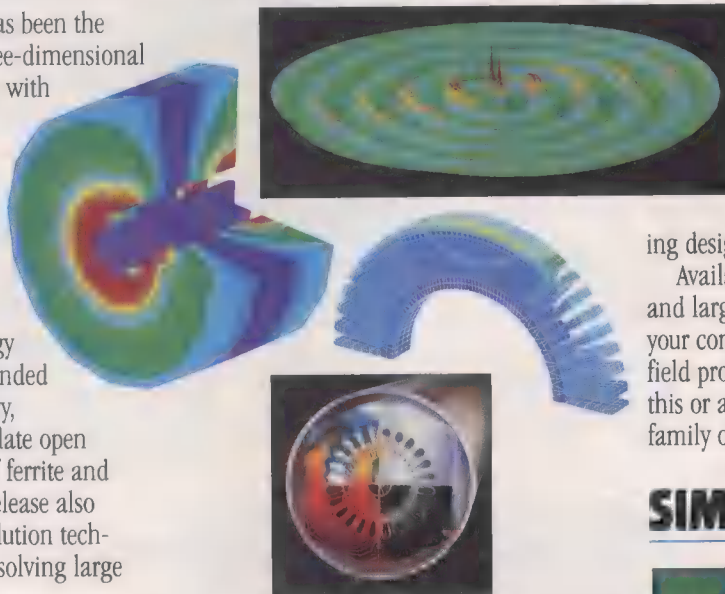
One of the first initiatives of Hale and Noyes was recruiting the physicist they regarded as the best in the United States: Robert Millikan, who was then at the University of Chicago. That romancing was interrupted by World War I, during which all three spent much of their time in Washington, D.C., with the newly formed National Research Council, created primarily through

(Continued on p. 16)

## A major advance in computational electromagnetics: MSC/EMAS Version 2.0

Since 1989, MSC/EMAS has been the benchmark for two- and three-dimensional finite element analysis. Now, with the release of MSC/EMAS Version 2.0, The MacNeal-Schwendler Corporation has raised the standard in electromagnetic field analysis even higher.

Version 2.0 introduces more leading edge technology with the addition of an expanded materials and element library, including the ability to simulate open boundaries and the effects of ferrite and other lossy materials. This release also features greatly enhanced solution techniques capable of efficiently solving large three-dimensional problems.



If you're involved in EMI/EMC, RF shielding, microwave heating, magnetic signature, antennas or other radiating structures, Version 2.0 is the ideal tool for reducing product development costs and improving design insight.

Available on most popular workstations and larger computers, MSC/EMAS can solve your company's toughest electromagnetic field problems. For a video or brochure on this or any member of MSC's complete family of software, call us at 414-357-8723.

### SIMPLY POWERFUL.



The  
MacNeal-Schwendler  
Corporation


# MSC/EMAS

MSC and MSC/ are registered service and trademarks of The MacNeal-Schwendler Corporation. MSC/EMAS is a trademark of MSC. NASTRAN is a registered trademark of NASA. MSC/NASTRAN is an enhanced proprietary solution developed, maintained, supported and marketed by MSC.

The Perfect Complement to MSC/NASTRAN



# Think of it as a mathematical playground for your mind.



## Open up a whole new world of scientific and engineering problem solving with Mathcad 3.1.

Between the challenge of new problems and the reward of new solutions there is math. Often, a lot of math.

Which is why we make Mathcad.

It's the remarkable software that gives you the power to perform everything from simple arithmetic to advanced computational gymnastics more quickly, easily and naturally than with calculators, spreadsheets or programming alone.

You use Mathcad just like a scratch

pad. Simply type your work anywhere in the live screen and Mathcad does the calculating. Change variables and Mathcad automatically updates the answers. It plots in 2-D and 3-D. And prints your results, complete with text and graphics, in presentation quality documents.

Plus you see your work more clearly because Mathcad uses the same math notation you would use with a pencil and paper.

### It's the complete math system.

Mathcad comes with more than 200 commonly used functions that let you quickly build even complex equations and formulas

with exponentials, differentials, cubic splines, FFTs and more.

Full symbolic calculation capabilities are available with a menu pick. So you can evaluate any integral, Taylor series or infinite sum just by clicking.

Optional Electronic Handbooks\* give you instant access to data, formulas and diagrams once found only in reference books. They're fully interactive. Just click and paste them right into your document and work with them like any other text, value, table or formula. Or work with them right in the handbook itself.

Plus optional Applications Packs with modifiable templates are available for all major engineering and science fields.

### Over 200,000 satisfied users.

Each day, more engineers, scientists and other technical professionals discover that Mathcad opens up a whole new world of freedom in problem solving, and helps them manage their time more efficiently.

Which is why in addition to being the most popular math software you can buy, it is also the most productive.

- Microsoft® Windows,™ Motif and Open Look interfaces make Mathcad easy to learn and use
- Optional Electronic Handbooks include The CRC Materials Science and Engineering Handbook,

NEW  
**UNIX 3.1**  
AVAILABLE

Machine Design and Analysis from the *Standard Handbook of Engineering Calculations* (McGraw-Hill), and the Mathcad Treasury of Methods and Formulas

- Optional Signal Processing Function Pack contains over 60 functions including FIR and IIR filter coefficients and power spectrum
- Optional Applications Packs with adaptable templates for Electrical, Mechanical, Civil and Chemical Engineering, Statistics, Advanced Math, and Numerical Methods
- PC Windows, PC DOS, UNIX® and Macintosh® versions available

### Free demo disk.

For a free Mathcad demo disk, or upgrade information, call 1-800-628-4223, (or 617-577-1017, Fax: 617-577-8829).

Or see your software dealer.

\*Electronic Handbooks require Mathcad 3.1.  
© 1992 MathSoft, Inc. TM and ® signify manufacturer's trademark or registered trademark respectively.



**1-800-MATHCAD**

**The answer is  
Mathcad®**

**MathSoft, Inc.** 201 Broadway, Cambridge, MA 02139 USA • Phone: 1-800-628-4223 • 617-577-1017 • Fax: 617-577-8829

Australia: Hearne (03) 866 1766; Belgium: SCIA 013/55 1775; Denmark: Engberg 042-251777; Finland: Zenex 0-692-7677; France: Uniware 01-45-272061; Germany: Softline (0 78 02) 4036; Italy: Channel 02-90091773; Japan: SMI 03-5476-9816; Sweden: AkademiData (018)2400 35; Switzerland: Redacom 032 41 01 11; UK: Adept Scientific (0462) 480055. In other locations, contact MathSoft, USA.



## Calendar

(Continued from p. 8)

more Convention Center, Baltimore, Md.; Doris Thomas, International Test Conference, 514 E. Pleasant Valley Blvd., Suite 3, Altoona, Pa. 16602; 814-941-4666.

**Petroleum and Chemical Industry Technical Conference (IA)**; Sept. 28-30; River City Marriott Hotel, San Antonio, Texas; Knox Pitzer, Thermon Manufacturing Co., 100 Thermon Dr., Box 609, San Marcos, Texas 78666; 512-396-5801.

**13th International Electronics Manufacturing Technology Symposium (CHMT)**; Sept. 28-30; Hyatt Regency Inner Harbor Hotel, Baltimore, Md.; Bill Moody, 2529 Eaton Rd., Wilmington, Del. 19810; 302-478-4143; fax, 302-478-7057.

**First IEEE International Conference on Universal Personal Communications (COM, Dallas Section)**; Sept. 29-Oct. 1; Loew's Anatole Hotel, Dallas; Dhawal Moghe, Bell Northern Research, 1150 E. Arapaho Rd., Richardson, Texas 75081; 214-997-4506; fax, 214-997-4792.

**Advanced Semiconductor Manufacturing Conference and Workshop (ED)**; Sept. 30-Oct. 1; Cambridge Hyatt Regency Hotel, Cambridge, Mass.; Margaret Bachmeyer, SEMI, 2000 L St., N.W., Suite 200, Washington, D.C. 20036; 202-457-9584; fax, 202-659-8534.

**Challenges in Optoelectronic Packaging (LEO, CHMT)**; Sept. 30-Oct. 1; Hyatt Regency Hotel, Baltimore, Md.; IEEE/LEOS, 445 Hoes Lane, Box 1331, Piscataway, N.J. 08855-1331; 908-562-3893; fax, 908-562-1571.

**Fourth Annual IEEE International Workshop on Computer Aided Modeling Analysis and Design of Communication Links and Networks (COM)**; Sept. 30-Oct. 2; Le Chateau Montebello, Montebello, Que., Canada; Hussein T. Mouftah, Department of Electrical Engineering, Queen's University, Kingston, Ont., K7L 3N6, Canada; 613-545-2934; fax, 613-545-6615.

**International Professional Communication Conference—IPCC '92 (PC)**; Sept. 30-Oct. 2; La Fonda on the Plaza Hotel, Santa Fe, N.M.; Susan Dressel, Information Services, Los Alamos National Laboratory, Mail Stop M704, Los Alamos, N.M. 87545; 505-667-6101; fax, 505-667-1754.

**International Workshop on Hardware-Software Codesign (C, CAS)**; Sept. 30-Oct. 2; Holiday Inn, Estes Park, Colo.; IEEE Computer Society, Conference Department, 1730 Massachusetts Ave., N.W., Washington, D.C. 20036-1903; 202-371-1013; fax, 202-728-0884.

### OCTOBER

**International Workshop on Intelligent Manufacturing Systems (CS et al.)**; Oct. 1-2; Hyatt Regency Hotel, Dearborn, Mich.; N.A. Kheir, Oakland University, Dodge Hall of Engineering, Rochester, Mich. 48309-4401; 313-370-2245; fax, 313-370-4261.

IEEE members attend more than 5000 IEEE professional meetings, conferences, and conventions held throughout the world each year. For more information on any meeting in this guide, write or call the listed meeting contact. Information is also available from: Conference Services Department, IEEE Service Center, 445 Hoes Lane, Box 1331, Piscataway, N.J. 08855; 908-562-3878; submit conferences for listing to: Ramona Foster, *IEEE Spectrum*, 345 E. 47th St., New York, N.Y. 10017; 212-705-7305.

For additional information on hotels, conference centers, and travel services, see the Reader Service Card.

## THE ULTIMATE CARD TRICK



## 320 MFLOPS

SuperCard i860-based vector processors deliver. With our Quad-860 multiprocessor and our SC-3XL ultimate performance models, CSPI continues to offer you the greatest range of vector processing power. Look what you can get on a 6U board with VME64:

- 320 MFLOPS Performance
- 64 MBytes Memory
- 200 MBytes per second I/O
- Intel's 50 MHz i860 XP chips
- VME64 and VSB
- MIL Spec
- GFLOPS Multiprocessors
- 200 usec 1024 complex FFT

**CSPI's SuperKit and Integration Toolkit software cut development time to the bone with:**

- Industry-leading Libraries
- pSOS+ Real-time Kernel
- UNIX compatibility
- DataFlo Visual Programming
- C, Fortran, and Ada

Call or write CSPI and learn more about our new Quad-860 and SC-3XL SuperCard family members: CSPI, 40 Linnell Circle, Billerica, MA 01821, **1-800-325-3110**, Fax: 508-663-0150. In Europe, CSPI, Creteil, France, 33-1-43774577, Fax: 33-1-43994508.



**CSPI**




**FREE  
DEMO PACKAGE!**

# IT'S HARD TO BELIEVE...



A SINGLE SOFTWARE PROGRAM – HiQ™ – CAN SOLVE PROBLEMS LIKE:

- *Analyzing the Deformation of a Nonlinear Elastic Beam*
- *Performing Linear and Nonlinear Circuit Analysis*
- *Modeling a Pumping Heart and Predicting Heart Failure*
- *Solving Electromagnetic Scattering Problems*
- *Computing the Trajectory of a Spacecraft*
- *Predicting the Levels of Smog Pollution*
- *Solving  Airline Scheduling Problem*

*Just ask HiQ customers like...*

Argonne National Labs  
Lockheed Corporation  
Dow Corning  
Magnavox  
National Institutes of Health

NASA  
Eastman Kodak  
Motorola  
General Electric  
Exxon Corporation



**Bimillennium**

First in Power Computing Software  
101 Albright Way, Los Gatos, CA 95030

**1-800-488-8662**

Circle No. 37

*Here's what they're saying about HiQ...*

*"HiQ is... Creating a Revolution for Engineers & Scientists!"* – Macintosh Engineering & Science Report



*"HiQ is in another league from the other computational programs on the market!"*

– Jeffrey Kane, M.D., Biomedical Researcher

*"It's a flexible environment... HiQ is head and shoulders above what we used in the past..."*


– Paul Toivonen, Sr. Project Eng., McDonnell Aircraft

HiQ, the HiQ logo, HiQ-Script, Problem Solvers and Bimillennium are trademarks of Bimillennium Corporation. Macintosh is a registered trademark of Apple Computer, Inc. Space Shuttle & Lunar photographs used with permission from NASA.



**MAC VERSION  
AVAILABLE NOW!  
UNIX VERSION  
FALL '92**







# 10ns BiCMOS 1 Megs from Motorola. Everything else is dead in the water.

1 Meg BiCMOS Fast SRAMs from Motorola demonstrate a simple evolutionary principle: survival of the fastest.

With 10ns access times at 1 Meg densities, nothing else even comes close enough to compare — for speed *and* density.

MCM6726	MCM6728	MCM6729 <sup>■</sup>	MCM67282 <sup>▲</sup>	MCM6727
128K x 8 bit	256K x 4 bit	256K x 4 bit	256K x 4 bit	1 meg x 1 bit
10 <sup>●</sup> , 12, 15ns	10, 12, 15ns	10, 12, 15ns	10, 12, 15ns	10, 12, 15ns
MCM6706A	MCM6705A	MCM6708A	MCM6709A <sup>■</sup>	MCM67082A <sup>▲</sup>
32K x 8 bit	32K x 9 bit	64K x 4 bit	64K x 4 bit	64K x 4 bit
8, 10, 12ns	10, 12ns	8, 10, 12ns	8, 10, 12ns	10, 12ns

● 3Q92 ■ Output Enable ■ Separate I/O

And as if that weren't enough to scare off the competition, these 1 Meg Fast SRAMs support both TTL and ECL I/O. They also feature an advanced pinout, with power supply, ground, and I/O pins centered on the package for reduced inductance and improved ground and power bussing.

Looking for even more speed? How about 8ns? That's the access time on our 256K BiCMOS Fast SRAMs.

Choose whichever speed-and-density combination is



right for you. Either way you'll get the built-in quality and reliability of Motorola's high volume, sub-micron manufacturing.

Reel in the power of our BiCMOS Fast SRAMs for your next design, and get ready to throw everything else back in the water.

To request technical information or a sample device, just mail in the coupon or FAX it to Motorola's Fast SRAM FAX line at 1-800-347-MOTO (6686).

Let's make some waves, Motorola.		SPEC 9/92	
Send me more on BiCMOS Fast SRAMs today.			
Motorola, Inc. Fast SRAM Operations, P.O. Box 1466, Austin, Texas 78767			
Name _____		Title _____	
Company _____			
Address _____			
City _____		State _____ Zip _____	
Phone _____			
Application _____			
Production Start Date _____		Estimated Usage: 1992 _____ 1993 _____	
SRAM: 1 Meg TTL I/O		1 Meg ECL I/O	
<input type="checkbox"/> 1 meg x 1 <input type="checkbox"/> 10ns <input type="checkbox"/> 256K x 4 <input type="checkbox"/> 12ns <input type="checkbox"/> 128K x 8 <input type="checkbox"/> 15ns		<input type="checkbox"/> 1 meg x 1 <input type="checkbox"/> 10ns <input type="checkbox"/> 256K x 4 <input type="checkbox"/> 12ns <input type="checkbox"/> 15ns	
		256K TTL I/O	
		<input type="checkbox"/> 64K x 4 <input type="checkbox"/> 8ns <input type="checkbox"/> 32K x 8 <input type="checkbox"/> 10ns <input type="checkbox"/> 32K x 9* <input type="checkbox"/> 12ns	
		*10 & 12ns Only	

If you like what's new, wait 'til you see what's next.



Motorola offers a complete portfolio of BiCMOS and CMOS Fast SRAMs with densities from 16K to 1 meg, plus 2 and 8 meg modules. CMOS access times are as fast as 15ns (256K) and 20ns (1 meg).

▲ and Motorola are registered trademarks of Motorola, Inc. © 1992 Motorola, Inc.




Now Available  
from Electros witch

## A Full Line of OPEN FRAME Rotary Switches

The Electros witch unit in Raleigh, NC is now manufacturing, engineering and selling the former Centra-lab line of open frame electronic rotary switches.

They will continue to be available through existing local distributors.

Four basic product groups: subminiature, miniature, heavy-duty and power configurations.



**ELECTROSWITCH**  
UNIT OF ELECTRO SWITCH CORP.

Raleigh, NC 27604

Telephone: (919) 833-0707 • Fax: (919) 833-8016

Circle No. 39

### HARMONICS or PHASE COUPLINGS

With Integrated Signal Analysis (ISA) you can considerably reduce the time it takes you to productively utilize recent advances in higher-order spectral analysis in analyzing, interpreting and modeling signal characteristics of nonlinear systems and phenomena such as harmonic generations or intermodulations.

ISA only needs your experimental data to provide the comprehensive linear and higher-order spectral analysis results.

The publication quality graphics are completely built-in. The textbook style Reference Guide explains the principles and concepts of linear and higher-order spectral analysis underlying ISA.


Workstation versions are also available for SUN SPARC and DEC VAX/VMS computers.

#### SOME APPLICATION AREAS

- Underwater Acoustics
- Nonlinear Modal Analysis
- Offshore Technology
- Nonlinear Ship Dynamics
- Nonlinear Radar
- Biomedical Research
- Nonlinear Vibrations
- Nonlinear Signatures
- Communications
- Oceanography
- Fluid Dynamics
- Plasma Physics

#### FEATURES

- Bispectrum, Bicoherency, Bicorrelation
- Quadratic Transfer Function
- Quadratic System Coherency
- Wavenumber-Frequency
- Linear Transfer Function
- FFT (prime factor based)
- Preconditioning Digital Filters
- Phase-Sensitive Complex Demodulation
- Standard/MEM Spectrum & Correlation



#### INTEGRAL SIGNAL PROCESSING, INC.

7719 Wood Hollow Drive, Suite 219, Austin, TX 78731-1634  
1-800-366-9081, Voice: (512) 346-1451, Fax: (512) 346-8290

Circle No. 43

## Calendar

(Continued from p. 14)

**GaAs Reliability Workshop (ED)**; Oct. 4; Fontainebleau Hilton Hotel, Miami Beach, Fla.; Anthony Immorlica, General Electric Co., Electronics Laboratory, Electronics Park, Syracuse, N.Y. 13221; 315-456-3514; fax, 315-456-0695.

**International Symposium on Time-Frequency and Time-Scale Analysis (SP)**; Oct. 4-6; Victoria Conference Center, Victoria, B.C., Canada; Jan Kvamme, Engineering Continuing Education, University of Washington, 4725 30th Ave., N.E., Seattle, Wash. 98105; 206-543-5539; fax, 206-543-2352.

**GaAs Integrated Circuits Symposium (ED)**; Oct. 4-7; Fontainebleau Hilton Hotel, Miami Beach, Fla.; Suzanne Kuntz, Courtesy Associates, 655 15th St., N.W., Suite 300, Washington, D.C. 20005; 202-347-5900; fax, 202-347-6109.

**International Telecommunications Energy Conference (PEL, COM)**; Oct. 4-8; J.W. Marriott Hotel, Washington, D.C.; Pete Paradissis, Reliance Telecommunications, 1122 F St., Lorain, Ohio 44052; 216-288-1122.

**Bipolar/BiCMOS Circuits and Technology Meeting (ED)**; Oct. 5-6; Marriott City Center Hotel, Minneapolis, Minn.; John Shier, VTC Inc., 2800 East Old Shakopee, Bloomington, Minn. 55425; 612-853-3292; fax, 612-853-3355.

**11th Symposium on Reliable Distributed Systems (C)**; Oct. 5-7; Windham Warwick Hotel, Houston, Texas; IEEE Computer Society, Conference Department, 1730 Massachusetts Ave., N.W., Washington, D.C. 20036-1903; 202-371-1013; fax, 202-728-0884.

**Digital Avionics Systems Conference (AES, Seattle)**; Oct. 5-8; Westin Hotel, Seattle, Wash.; Jose Bolanos, Boeing, M/S 96-16, Box 3707, Seattle, Wash. 98124; 206-237-3719; fax, 206-237-6088.

**International SOI Conference (ED)**; Oct. 6-8; Marriott at Sawgrass Resort, Ponte Vedra, Fla.; Jerry Brandewie, Sematech (Rockwell International), 2706 Monopolis Dr., Austin, Texas 78741; 512-356-3449; fax, 512-356-3521.

**International Conference on Computer Design: VLSI in Computers and Processors—ICCD '92 (ED)**; Oct. 11-14; Royal Sonesta Hotel, Cambridge, Mass.; IEEE Computer Society, 1730 Massachusetts Ave., N.W., Washington, D.C. 20036-1903; 202-371-1013; fax, 202-728-0884.

**Military Communications Conference (COM)**; Oct. 11-14; Sheraton Harbour Island Hotel, San Diego, Calif.; John Peckham, General Dynamics Corp., Box 85468, San Diego, Calif. 92138; 619-573-5452; fax, 619-592-5320.

**International Symposium on Systems, Man and Cybernetics (SMC)**; Oct. 18-21; Knickerbocker Hotel, Chicago; Richard Saeks, Department of Electrical and Computer Engineering, Illinois Institute of Technology, Chicago, Ill. 60616; 312-567-3221.


**VHDL International Users Forum (COMP)**; Oct. 18-21; Omni Shoreham Hotel, Washington, D.C.; VHDL International Inc., 407 Chester St., Menlo Park, Calif. 94025; 800-554-2550, 415-329-0578.

**Sixth Annual Leesburg Workshop on Concurrent Engineering (R)**; Oct. 19-22; Xerox Training Center, Leesburg, Va.; Henry N. Hartt, Vitro Corp., Suite 300, West Wing, 600 Maryland Ave., S.W., Washington, D.C. 20024; 202-646-6339; fax, 202-646-6398.

**Visualization '92 (C)**; Oct. 19-23; Boston Park Plaza Hotel, Boston; IEEE Computer Society, Conference Department, 1730 Massachusetts Ave., N.W., Washington, D.C. 20036-1903; 202-371-1013.

**Workshop on Power Electronics in Transportation**  
(Continued on p. 14H)





## Introducing the DAS-1600. Fasten your seatbelts.

Here's a DAS-1600 "Fast Track" special to help you get started.

With all new EASYEST AG™ software from Keithley Asyst when you buy an \$899

DAS-1600 board by December 31, 1992, the \$899 offer is available.

Call 800-348-0033

now and only for Extra Value to tablet or for more information.

MS or VISA accepted

FAX: 508-880-0179

I want supercharged  
Data Acquisition performance.

**\$899**

- ☐ Please send me more details on the DAS-1600.
- ☐ Please send me a free EASYEST AG software demo disk.
- ☐ Please have a sales representative call.

Phone \_\_\_\_\_

Name \_\_\_\_\_

Title \_\_\_\_\_

Company \_\_\_\_\_

Address \_\_\_\_\_

City \_\_\_\_\_ State \_\_\_\_\_

Zip/Postal Code \_\_\_\_\_ Country \_\_\_\_\_

**KEITHLEY METRABYTE**  
DATA ACQUISITION

DAS1A

for greater flexibility.  
anced ASIC design  
wer cost and higher  
ility. □ High-quality  
rinted circuit board  
Pop Up Control Panel  
Optional Windows®  
sive third-party soft-

**"Fast Track" offer.**

**ew EASYEST AG™**

**iley Asyst when**

**10 board by**

**quick start-up,**

**osition to disk, and**

**(normally \$395).**

**and an offer like this,**

**even faster.**

**Call 800-348-0033**

and ask for Cora Vette to order or for more information. And give yourself the winning edge.

□ Burst mode to simulate simultaneous sample and hold. □ 8 differential or 16 single-ended analog inputs. □ 32 bits of digital I/O. □ 2, 12-bit D/As with 5 selectable

**KEITHLEY METRABYTE**  
DATA ACQUISITION

IEE992 DAS1

Keithley Data Acquisition, 440 Myles Standish Blvd., Taunton, MA 02780, FAX: 508-880-0179  
©1992 Keithley Instruments, Inc. Keithley Metrabyte and Keithley Asyst are registered trademarks and EASYEST AG is a trademark of Keithley Instruments, Inc.  
Windows is a registered trademark of Microsoft Corporation.




**Now Available  
from Electros witch**

**A Full Line of  
OPEN FRAME  
Rotary Switches**

The Electros witch unit in Raleigh, NC is now manufacturing, engineering and selling the former Centra-lab line of open frame electronic rotary switches.

They will continue to be available through existing local distributors.

Four basic product groups: subminiature, miniature, heavy-duty and power configurations.



UNIT OF ELECTRO SWITCH CORP.

Raleigh, NC 27601

Telephone: (919) 833-0707 • Fax: (919) 833-0708

Circle No. 39

**HARMONICS or PHASE COUPLING**

With Integrated Signal Analysis (ISA) you can reduce the time it takes you to productively utilize *higher-order spectral analysis* in analyzing, in signal characteristics of nonlinear systems and phenomena or intermodulations.

ISA only needs your experimental data to provide *higher-order spectral analysis* results.

The publication quality graphics are completely built in. The *Harmonics or Phase Coupling* Guide explains the principles and concepts of *spectral analysis* underlying ISA.


Workstation versions are also available for SUN SPARC workstations.

**SOME APPLICATION AREAS**

- Underwater Acoustics
- Nonlinear Modal Analysis
- Offshore Technology
- Nonlinear Ship Dynamics
- Nonlinear Radar
- Biomedical Research
- Nonlinear Vibration
- Nonlinear Signal Processing
- Communication Systems

**FEATURES**

- Bispectrum, Bicoherency, Bicoherence
- Quadratic Transfer Function
- Quadratic System Coherency
- Wavenumber-Frequency
- Linear Transfer Function
- FFT (prime factor based)
- Preconditioning Digital Filters
- Phase-Sensitive Complex Demodulation
- Standard/MEM Spectrum & Correlation



**INTEGRAL SIGNAL PROCESSING, INC.**

7719 Wood Hollow Drive, Suite 219, Austin, TX 78731-1634  
1-800-366-9081, Voice: (512) 346-1451, Fax: (512) 346-8290

Circle No. 43

**Calendar**

(Continued from p. 14)

**GaAs Reliability Workshop (ED)**; Oct. 4; Fontainebleau Hilton Hotel, Miami Beach, Fla.; Anthony Immorlica, General Electric Co., Electronics Laboratory, Electronics Park, Syracuse, N.Y. 13221; 315-456-3514; fax, 315-456-0695.

**International Symposium on Time-Frequency and Time-Scale Analysis (SP)**; Oct. 4-6; Victoria Conference Center, Victoria, B.C., Canada; Jan Kvamme, Engineering Continuing Education, University of Washington, 4725 30th Ave., N.E., Seattle, Wash. 98105; 206-543-5539; fax, 206-543-2352.

**GaAs Integrated Circuits Symposium (ED)**; Oct. 4-7; Fontainebleau Hilton Hotel, Miami Beach, Fla.; Suzanne Kuntz, Courtesy Associates, 655 15th St. N.W., Suite 200,

**International SOI Conference (ED)**; Oct. 6-8; Marriott at Sawgrass Resort, Ponte Vedra, Fla.; Jerry Brandewie, Sematech (Rockwell International), 2706 Monopolis Dr., Austin, Texas 78741; 512-356-3449; fax, 512-356-3521.

**International Conference on Computer Design: VLSI in Computers and Processors—ICCD '92 (ED)**; Oct. 11-14; Royal Sonesta Hotel, Cambridge, Mass.; IEEE Computer Society, 1730 Massachusetts Ave., N.W., Washington, D.C. 20036-1903; 202-371-1013; fax, 202-728-0884.

**Military Communications Conference (COM)**; Oct. 11-14; Sheraton Harbour Island Hotel, San Diego, Calif.; John Peckham, General Dynamics Corp., Box 85468, San Diego, Calif. 92138; 619-573-5452; fax, 619-592-5320.

**International Symposium**



NO POSTAGE  
NECESSARY  
IF MAILED  
IN THE  
UNITED STATES

**BUSINESS REPLY MAIL**

FIRST CLASS MAIL PERMIT NO. 72 TAUNTON, MA

POSTAGE WILL BE PAID BY ADDRESSEE

**KEITHLEY METRABYTE**  
DATA ACQUISITION

440 Myles Standish Blvd.  
Taunton, MA 02780-9962




Department, 1730 Massachusetts Ave., N.W., Washington, D.C. 20036-1903; 202-371-1013; fax, 202-728-0884.

**Digital Avionics Systems Conference (AES, Seattle)**; Oct. 5-8; Westin Hotel, Seattle, Wash.; Jose Bolanos, Boeing, M/S 96-16, Box 3707, Seattle, Wash. 98124; 206-237-3719; fax, 206-237-6088.

**Visualization '92 (C)**; Oct. 19-23; Boston Park Plaza Hotel, Boston; IEEE Computer Society, Conference Department, 1730 Massachusetts Ave., N.W., Washington, D.C. 20036-1903; 202-371-1013.

**Workshop on Power Electronics in Transportation**  
(Continued on p. 14H)





## Introducing the DAS-1600. Fasten your seatbelts.

**Turbocharged data acquisition performance at an economy price—only \$899.**

When we introduced the DAS-16 in the early '80s, it raced quickly into the top spot among general purpose data acquisition boards.

Well, here we go again. We've just taken the wraps off the next generation: the DAS-1600. A board with breakthrough performance and the best software support in the industry. For the incredible price of \$899—less than our original DAS-16.

Packed with advanced features you need for the '90s. Including:

- 100 ksample/sec with 12-bit resolution.
- Burst mode to simulate simultaneous sample and hold.
- 8 differential or 16 single-ended analog inputs.
- 32 bits of digital I/O.
- 2, 12-bit D/As with 5 selectable



ranges for greater flexibility.

- Advanced ASIC design for lower cost and higher reliability.
- High-quality

four-layer printed circuit board

for reduced noise. □ Pop Up Control Panel for quick start-up. □ Optional Windows® 3.1 software. □ Extensive third-party software support.



### **Special "Fast Track" offer.**

**\$300 off new EASYEST AG™ software from Keithley Asyst when you buy a DAS-1600 board by**

**Oct. 31, 1992.** Get quick start-up, troubleshooting, acquisition to disk, and graphics for only \$95 (normally \$395). With a price of \$899, and an offer like this, the DAS-1600 will go even faster.

**Call 800-348-0033**

and ask for Cora Vette to order or for more information. And give yourself the winning edge.

**KEITHLEY METRABYTE**  
DATA ACQUISITION

IEE992 DAS1



## Announcing PDE2D

### a Descendant of PDE/Protran

A new product which solves quite general time-dependent, steady-state and eigenvalue systems of partial differential equations, in general two-dimensional regions, is now offered by Granville Sewell, author of PDE/PROTRAN. The new program, **PDE2D**, features an interactive user interface, and can solve nonlinear problems of the form:

$$C(X,Y,T,U) \frac{\partial U}{\partial T} = \frac{\partial}{\partial X} A(X,Y,T,U,U_x,U_y) + \frac{\partial}{\partial Y} B(X,Y,T,U,U_x,U_y) - F(X,Y,T,U,U_x,U_y)$$

or

$$\frac{\partial}{\partial X} A(X,Y,U,U_x,U_y) + \frac{\partial}{\partial Y} B(X,Y,U,U_x,U_y) = F(X,Y,U,U_x,U_y)$$

where C, A, B, F, and U may be vector functions. **PDE2D** also solves the corresponding eigenvalue problem, and quite general boundary conditions can be handled.

**PDE2D** offers a choice of quadratic, cubic or quartic isoparametric triangular elements, and the resulting large linear systems may be solved by a band solver, a frontal method, or an iterative bi-conjugate gradient method. Automatic refinement and grading of the triangular mesh are available, and extensive graphical output capabilities are provided.

Readers interested in **PDE2D** can obtain the interactive driver free of charge, for evaluation purposes, from Dr. Sewell, at the address given below.

Granville Sewell  
P.O. Box 12141  
El Paso, TX 79913  
Telephone: (915) 581-3763  
e-mail: xxvd004@math.ep.utexas.edu

Circle No. 38



### A Workshop on Wavelets and their Applications October 22 - 23, 1992

Aware, Inc.  
One Memorial Drive  
Cambridge, MA 02142

The staff of Aware Inc and affiliated academic consultants will conduct a workshop on Wavelets as a mathematical structure and their practical application to several areas of current interest. These will include

- Lossy Compression
- Channel Coding
- Numerical Solution of Partial Differential Equations

The workshop will explore the advantages and disadvantages of using wavelets for these and other applications. Attendance will be limited to 20 persons to allow for interactive discussion. A tuition of \$1000 will be charged. This will include the cost of lunch, coffee breaks and reception. Attendees will receive lecture notes of the several presentations, a videotape illustrating the material, and a software package introducing the basic Wavelet constructions.

For information, contact Christina Gorecki at  
AWARE, Inc.  
One Memorial Drive  
Cambridge, MA 02142

Phone: 617-577-1700 Fax: 617-577-1710  
Email: workshop@aware.com  
We accept Mastercard and Visa

Circle No. 30

## Calendar

(Continued from p. 14D)

(PEL); Oct. 22-23; Hyatt Regency Hotel, Dearborn, Mich.; V. Anand Sankaran, Ford Motor Co., Room S-2037, Scientific Research Laboratory, 20 000 Rotunda Dr., Dearborn, Mich. 48121-2053; 313-390-8689.

**Wafer Level Reliability Workshop (ED)**; Oct. 25-28; Stanford Sierra Lodge, Lake Tahoe, Calif.; Harry Schaft, National Institute of Standards and Technology, Building 225, Room B360, Route 270, Quince Orchard Road, Gaithersburg, Md. 20899; 301-975-2234.

**26th Annual Asilomar Conference on Signals, Systems, and Computers (SP, C)**; Oct. 26-28; Asilomar Hotel, Pacific Grove, Calif.; James A. Ritcey, Department of Electrical Engineering, FT-10, University of Washington, Seattle, Wash. 98195; 206-543-4702.

### NOVEMBER

**Regional Symposium on Electromagnetic Compatibility (EMC, Region 8)**; Nov. 2-5; Tel Aviv Hilton Hotel, Israel; Symposium Secretariat, Ortra Ltd., Box 50432, Tel Aviv 61500, Israel; Dani Tider, (972+3) 664 825.

**International Conference on Computer-Aided Design (ED)**; Nov. 9-12; Santa Clara Convention Center, Santa Clara, Calif.; IEEE Computer Society, 1730 Massachusetts Ave., Washington, D.C. 20036-1903; 202-371-1013.

**Technologies Enabling Tomorrow (C et al.)**; Nov. 11-13; World Congress Center, Melbourne, Australia; Marg Scarlett, Convention Network, 224 Rouse St., Port Melbourne, Victoria 3207, Australia; (61+3) 646 4122; fax, (61+3) 646 7737.

**Second Workshop on the Management of Replicated Data (C)**; Nov. 12-13; Monterey Marriott Hotel, Monterey, Calif.; Jehan-François Paris, Department of Computer Sci-

ence, University of Houston, Houston, Texas 77204-3475; 713-749-3943.

**New Generation Knowledge Engineering (C)**; Nov. 16-18; Doubletree Washington National Airport Hotel; Julie Walker, 11820 Parklawn Dr., Rockville, Md. 20852-2529.

**LEOS '92 Annual Meeting (LEO)**; Nov. 16-19; Hynes Convention Center, Boston; IEEE/LEOS, 445 Hoes Lane, Box 1331, Piscataway, N.J. 08855-1331; 908-562-3896.

**International Conference on Communication Systems/International Symposium on Information Theory and Its Applications—ICCS/ISITA '92 (IT, COM)**; Nov. 16-20; Westin Stamford and Westin Plaza Hotels, Singapore; Esther Yeo, Mansfield International, 71 Robinson Rd., 4th Storey, Crosby House, 0106, Singapore; (65) 224 0000.

**Wescon '92 (Region 6, Los Angeles Council)**; Nov. 17-19; Anaheim Convention Center, Anaheim, Calif.; Electronic Conventions Management, 8110 Airport Blvd., Los Angeles, Calif. 90045; 213-215-3976 or 800-877-2668.

**18th Annual Convention and Exhibition (ACE '92) (Calcutta, India C)**; Nov. 21-23; Birla Industrial and Technological Museum, Calcutta, India; Secretariat, c/o Department of E&T.C.E., Jadavpur University, Calcutta-700 032, India; (91+72) 2851; telex, (91) 0215195.

**First Asian Test Symposium (C)**; Nov. 26-27; Hiroshima Grand Hotel, Hiroshima, Japan; IEEE Computer Society, Conference Department, 1730 Massachusetts Ave., N.W., Washington, D.C. 20036-1903; 202-371-1013; fax, 202-728-0884.

### DECEMBER

**37th Annual Conference on Magnetism and Magnetic Materials (MAG)**; Dec. 1-4; Westin Galleria Hotel, Houston,


(Continued on p. 14J)



# 486-B<sup>2</sup>T

## 486/860 Speed... Microway Quality.


Microway has engineered four distinctive black tower systems. The 486-B<sup>2</sup>T is designed for high-end users. It comes standard with American 486 motherboards and power supplies, yet has a reasonable starting price of \$2,195. A broad range of options can be installed including high speed and capacity hard disks, intelligent serial controllers, tape back-up units, high end graphics adapters and our *Number Smasher-860*. These systems are ideal for configuring Novell or UNIX file servers, multiuser systems, and workstations for graphics, CAD and scientific uses. The 486-B<sup>2</sup>T comes with dual fans. *Across the Board*™ Cooling and American industrial grade power supplies. All systems are thoroughly tested, burned in and include the best technical support in the industry, which we've provided since 1982.



**Number Smasher-860**

The key to attaining workstation performance is Microway's 40MHz *Number Smasher-860*. It features a four-way interleaved 64-bit memory system that runs at 160 megabytes/sec. The *Number Smasher's* i860 has been clocked at 80 megaflops doing matrix multiplies, 67 megaflops doing FFTs and 11.8 Double Precision Linpack Megaflops on large arrays—ten times the speed of a 486 and twice the speed of a Cray 1F! One happy user recently reported that his "Baby Cray" was happily humming away saving him thousands of dollars per month in 3090 rentals. The *Number Smasher* comes with the finest i860 compilers on the market, your choice of Microway's NDP™ FORTRAN, C/C++ or Pascal.

**Call or write today for more information on Microway's new black tower systems.**



# Microway®

**Technology You Can Count On**

Corporate Headquarters, Box 79, Kingston, MA 02364 USA • TEL 508-746-7341 • FAX 508-746-4678 • U.K./Europe 081-541-5466  
France 01 43 2 69593 • Germany 069-75-2023 • Holland 40 836455 • Italy 02-74.90.749 • Japan 0474 23 1322 • Norway 6-892020

Circle No. 19



## Calendar

(Continued from p. 14H)

Texas; Diane Suiters, Courtesy Associates, 655 15th St., N.W., Suite 300, Washington, D.C. 20005; 202-639-5088.

**Global Telecommunications Conference (COM)**; Dec. 6-9; Buena Vista Palace Hotel, Orlando, Fla.; Ron Kandell, Siemens/Stromberg-Carlson, 900 Broken Sound Parkway, Boca Raton, Fla. 33487; 407-955-8230; fax, 407-955-8771.

**Asia-Pacific Conference on Circuits and Systems (CAS)**; Dec. 8-11; Hyatt Kingsgate Hotel, Sydney, Australia; Heather Harriman, IREE Executive Director, IREE Australia, Commercial Unit 3, 2 New McLean St., Box 79, Edgecliff, NSW 2027, Australia; (61+2) 327 4822.

**23rd Annual IEEE Semiconductor Interface Specialists Conference (ED)**; Dec. 9-12; San Diego Hilton Hotel, San Diego, Calif.; Lalita Manchanda, AT&T Bell Laboratories, Crawford Corners Road, M/S 4C 406, Holmdel, N.J. 07733; 908-949-1679.

**International Electron Devices Meeting (ED)**; Dec. 13-16; San Francisco Hilton & Towers Hotel, San Francisco; Melissa Widerkehr, Suite 610, 1545 18th St., N.W., Washington, D.C. 20036; 202-986-1137.

**Conference on Decision and Control (CS)**; Dec. 16-18; Westin La Paloma Resort/Hotel, Tucson, Ariz.; T. Basar, Coordinated Science Laboratory, University of Illinois, 1101 West Springfield Ave., Urbana, Ill. 61801; 217-333-3607; fax, 217-244-1764.

### 1993 FEBRUARY

**International Solid-State Circuits Conference (SSC et al.)**; Feb. 24-26; San Francisco Marriott, San Francisco, Calif.; Diane S. Suiters, 655 15th St., N.W., Suite 300, Washington, D.C. 20005; 202-639-4255; fax, 202-347-6109.

### MARCH

**Multi-Chip Module Conference (ED)**; March 16-19; Cocoanut Grove, Santa Cruz, Calif.; S. Simon Wong, CIS-202, Stanford University, Stanford, Calif. 94305-4070; 415-725-3706; fax, 415-725-6949.

**International Reliability Physics Symposium (ED)**; March 21-25; Hyatt Regency, Atlanta, Ga.; David A. Baglee, 5604 Cometa Court N.E., Albuquerque, N.M. 87111; 505-893-3446; fax, 505-893-1049.

**International Conference on Micro-electronic Test Structures (ED)**; March 22-25; Gran Sitges Hotel, Barcelona, Spain; Loren W. Linholm, National Institute of Standards and Technology, B360 Technology Bldg., Gaithersburg, Md. 20899; 301-975-2052; fax, 301-948-4081.

### MAY

**Custom Integrated Circuits Conference (ED)**; May 9-12; Town & Country Hotel, San Diego, Calif.; Roberta Kaspar, 1597 Ridge Rd. W., Suite 101C, Rochester, N.Y. 14615; 716-865-7164.

**Photovoltaic Specialists Conference (ED)**; May 10-14; Galt House, Louisville, Ky.; Eldon C. Boes, NREL, Suite 710, 409 12th St., S.W., Washington, D.C. 20024; 202-484-1090; fax, 202-484-8177.

**International Symposium on Power Semiconductor Devices and ICs (ED)**; May 17-19; Hyatt Regency Monterey, Monterey, Calif.; M. Ayman Shibib, AT&T Bell Laboratories, Box 13566, Reading, Pa. 19612-3566; 215-939-6576; fax, 215-939-6795.

Your Single Source for  
The **W-I-D-E-S-T** Range  
of Electronic and Power

## ROTARY SWITCHES

We can match your needs:

- APPLICATIONS
- RATINGS
- OPTIONS
- SIZES

Complete lines include

- Electronic
- Minature
- Heavy-Duty
- Custom Designs
- Snap, Detent and Cam Actions
- Military & Enhanced Commercial

Our switch ratings range from under 1 amp to 200 amps



**ELECTROSWITCH**



UNIT OF ELECTRO SWITCH CORP.

Weymouth, Massachusetts 02188

Telephone: 617/335/5200 • FAX: 617/335/4253

Circle No. 35





Electrostatic potentials for an imide. Silver spheres represent 0.011 eV oxygen potential. The surface represents the mapped potentials, deformed and colored by data values.

Now there's a supercomputer  
for the left brain...

...and the right.

The IBM POWER Visualization System. To the left brain, it's a very high speed, high throughput processor for handling very large data sets. It links 8 to 32 processors in parallel to provide a peak rating of 1,280 double-precision megaflops\*. And it can access hundreds of gigabytes at up to 100 megabytes/second.

To the right brain, it's a unique visualization environment that lets you interactively explore the meaning of your data. For example, you can create and play back true-color, 3-D animations with up to 1,920 x 1,536 pixels. You can easily manipulate or correlate data.

Zoom in on details at high resolution. Or even navigate through images in real time.

The IBM POWER Visualization System is the supercomputer designed from the ground up for high speed number-crunching integrated with full-featured visualization. And it's supported by tools for developing interactive simulations and computational steering applications.

Anyway you look at it, it deserves a closer look. For a free brochure, call (800) 388-9820. Or write: IBM Scientific Visualization Systems, T.J. Watson Research Center, P.O. Box 704, Yorktown Heights, NY 10598.

\*925 double-precision megaflops max, LINPACK TPP.

IBM is a registered trademark of International Business Machines Corporation. © 1992 IBM Corporation.





## Books

(Continued from p. 12)

the efforts of Hale. At the end of the war, all three returned to Pasadena, and in 1920 the institute took on its current name.

A new era had begun. As the author points out, the structure of scientific research in the United States was yet to be determined. Research in Europe was conducted mainly in free-standing institutes. The role of the National Research Council, and its relation to the Government and to scientific research in general, was not clear. Millikan was opposed to expanding the role of government in science, and after much soul-searching concluded that the most effective solution would be to combine the functions of research and instruction within the university environment. This was in effect the formulation of our modern research university. Thus began the campaign by this remarkable trio to embody this concept in their new Caltech.

This book details their efforts, including the perennial and always paramount problem of fund-raising. The role played by many generous and far-sighted benefactors was crucial in the development of the school. The account of the difficulties in financing a private university and of the vital importance of dedicated and skillful fund-raisers

is just as true today as it was in the 1920s.

The description of how the trio went about assembling an all-star faculty is fascinating, with lessons that are still useful. When I went to Caltech as a graduate student in 1941, I was lucky enough to meet many members of that original team. I did not fully appreciate then how privileged I was to have professors like Paul Epstein, Fritz Zwicky, Harry Bateman, and Theodor von Karman, in classes with just a handful of students. It was a unique environment.

*Millikan's School* is not just the definitive history of Caltech—it is a study of the issues facing science and engineering education. The account of the personal styles of each of the three—Hale, Millikan, and Noyes—is entertaining and instructive. I recommend this book not only to Caltech alumnae but to anyone interested in engineering, science, education, or their development in the 20th century United States.

*Allen E. Puckett is chairman emeritus of Hughes Aircraft Co., Los Angeles, which he joined immediately after earning a doctorate at Caltech in 1949. At Hughes Puckett headed a number of departments and laboratories before serving as chairman of the board and chief executive officer from 1978 to 1987. He has received many awards, including the IEEE's Frederik Philips Award in 1983 for managing the development of communications satellites.*

## Electrification and culture in Chicago

Ronald R. Kline

**The Electric City: Energy and the Growth of the Chicago Area, 1880-1930.** Platt, Harold L., University of Chicago Press, Chicago, 1991, 352 pp., \$34.95.



In this well-researched book, an urban historian sets the history of the electrification of a major city in the illuminating social, economic, and political contexts of the progressive period. Arguing against what he sees as the "technological determinism" of historian Thomas P. Hughes, the author, Harold L. Platt, contends that "contemporary social values and cultural orientations determined the pace and direction of technological change, . . . not the inexorable logic of the machine."


Platt concludes that electricity performed the dual role of a "tool" and an "active agent for cultural change." Yet readers will probably come away from this book with the dis-

(Continued on p. 78)

## INTERNATIONAL JOURNALS AND BOOKS INFORMATION

### CONTROL

THEORY AND ADVANCED TECHNOLOGY



#### Editors

**Yoshifumi Sunahara**  
Professor Emeritus  
Kyoto Institute of Technology  
Japan


**D.P. Atherton**  
Professor  
The University of Sussex  
England

\*Subscription rate for one year (four issues)  
US\$130.00 (surface mail postage included)  
US\$160.00 (airmail postage included)  
¥30,900 (postage included)

Circle No. 10

### OPTOELECTRONICS

Devices and Technologies




#### Editors

**Makoto Kikuchi**  
Executive Technical Adviser  
SONY Corporation  
Japan

**Yoshihiro Hamakawa**  
Professor  
Osaka University  
Japan

\*Subscription rate for one year (two issues)  
US\$65.00 (surface mail postage included)  
US\$75.00 (airmail postage included)  
¥10,300 (postage included)

Circle No. 11



**Recent Advances in Mathematical Theory of Systems, Control, Networks and Signal Processing I, II**

*Proceedings of the International Symposium MTNS-91*


\*Edited by H. Kimura and S. Kodama

\*[I] 676 pages; [II] 638 pages

\*each volume ¥6,000 / US\$60.00

plus postage at cost

Circle No. 12



**"What Should Be Done?" "What Can Be Done?"**

*The Proceedings of NISTEP the Second International Conference on Science and Technology Policy Research*

\*Edited by S. Okamura

K. Murakami

I. Nonaka

\*416 pages

\*¥18,000 / US\$160.00 plus postage at cost

Circle No. 13

Ochanomizu Center Bldg., 2-12, Hongo 3-chome, Bunkyo-ku, Tokyo 113, Japan

MITA PRESS

mita

Tel: 03(3818)1011  
Telex: MITAPS 2722813  
Fax: 03(3818)1016



# Mathematica

## THE DEFINITIVE SYSTEM FOR TECHNICAL COMPUTATION

**"Not merely a product but a revolution"**

— Macworld


**"The importance of the program cannot be overlooked"**

— New York Times


**Basic function:** Integrated environment for numerical, symbolic, graphical computation, interactive programming.

**Users:** Scientists, engineers, mathematicians, programmers, financial analysts, students. Over 150,000 worldwide. Includes all 50 largest U.S. universities.

**Numerical computation:** Arbitrary-precision arithmetic, complex numbers, special functions (hypergeometric, elliptic, etc.), combinatorial and integer functions. Matrix operations, root finding, function fitting, Fourier transforms, numerical integration, numerical solution of differential equations, function minimization, linear programming.




**Symbolic computation:** Equation solving, symbolic integration, differentiation, power series, limits. Algebraic operations, polynomial expansion, factorization, simplification. Operations on matrices, tensors, lists, strings.




**Graphics and sound:** 2D, 3D plots of functions, data, geometrical objects. Contour, density plots. 3D rendering with intersecting surfaces, lighting models, symbolic descriptions. Color PostScript output, combining and labeling, publication quality graphics, animation (most versions). Sound from waveforms and data (most versions).

**Programming:** High-level, interactive, symbolic language. Procedural and functional programming



constructs. Transformation rules and pattern matching. Fully compatible on all platforms. No built-in limitations on computation size.



**External interface:** Input of data (numbers, records, text) from files, programs. Output in TeX, C, Fortran, PostScript. Calling of external programs and functions. General MathLink interprocess communication mechanism.

**User interface:** Electronic book interactive documents mixing text, graphics, animations, calculations. Graphics, animation, sound interapplication compatibility. Style sheets, hierarchical outlining. Computation kernel can run on remote computer (most versions).

**Additional material:** Journal, newsletters, more than 30 books. Add-on packages, free MathSource electronic resource.

**Versions:** Macintosh • Microsoft Windows • MS-DOS • NEC PC • DEC RISC, VAX • HP • IBM RISC • MIPS • NeXT • SGI • Sony • SPARC, Sun-3 • CONVEX and others • Network licensing available. Student versions. Now shipping Version 2.1.

For the latest information call Wolfram Research at:

**1-800-441-MATH**



**Wolfram Research**

**Wolfram Research, Inc.**

100 Trade Center Drive, Champaign, IL 61820-7237, USA  
217-398-0700; fax: 217-398-0747; email: info@wri.com

**Wolfram Research (UK) Ltd.**

Evenlode Court, Main Road, Long Hanborough, Oxon OX8 2LA, UK  
+44-(0)993-883400; fax: +44-(0)993-883800; email: info-uk@wri.com



## IEEE 488.2

for Your PC



### PC, PC AT, and PS/2 Boards

- 100% IEEE 488.2 compatible
- NAT4882™ IEEE 488.2 Controller chip
  - Optimized GPIB functionality
- Turbo488® performance chip
  - Over 1 Mbytes/sec reads and writes\*
- Includes NI-488.2™ for DOS and Windows software

\* Between boards; adhering to IEEE 488.1 Section 5.2.3.

### Windows and DOS Software

- High-speed IEEE 488.2 routines
- Standard NI-488® functions
- HP-style commands
- Standard Windows DLL entry points
- Interactive control utility
- Controller and Talker/Listener functions

### Call for FREE 1993 Catalog

See us at Autotestcon, booth 605



The Software is the Instrument™

6504 Bridge Point Parkway

Austin, TX 78730-5039

Tel: (512) 794-0100

(800) 433-3488

(U.S. and Canada)

Fax: (512) 794-8411

#### Branch Offices

AUSTRALIA 03 879 9422 • BELGIUM 02 757 00 20

CANADA 519 622 9310 • DENMARK 45 76 73 22

FINLAND 90 524566 • FRANCE 1 48 65 33 70

GERMANY 089 714 50 93 • ITALY 02 48301892

JAPAN 03 3788 1921 • NETHERLANDS 01720 45761

NORWAY 03 846866 • SWEDEN 91 896 0675

SWEDEN 08 984970 • SWITZERLAND 056 27 00 20

U.K. 0635 523545

© Copyright 1992 National Instruments Corporation. All rights reserved. Product and company names listed are trademarks or trade names of their respective companies.

Circle No. 22

# The engineer at large

## Salaries much too stable

Salary increases for professional engineers were hurt by the recession last year, with the median salary gain only 2.2 percent, according to the "1992 Income and Salary Survey" of members of the National Society of Professional Engineers (NSPE), Washington, D.C.

The median total annual income of respondents to the survey was US \$58 240, up from \$57 000 the previous year. These figures include income from their primary jobs, bonuses, and commissions.

The 2.2 percent gain is down from the 2.7 percent increase of the year before. However, the consumer price index rose by only 2.6 percent in 1991 whereas the year before the index went up 5.7 percent. This might indicate that the relative situation of engineers had improved.

The highest median income by major branch of engineering goes to the respondents in petroleum engineering (\$69 500). This group is followed by those in nuclear (\$65 000), chemical (\$64 000), mining (\$62 400), and sanitary, electronics, electrical, environmental, mechanical, and civil (primarily structural) engineering (all in the \$58 000 to \$62 109 range).

At the low end of the income spectrum are those employed in aeronautical and aerospace (\$56 746), industrial (\$56 000), civil (\$55 843), agricultural (\$51 900), and, last and not least, manufacturing engineering (\$51 185).

As in previous NSPE salary surveys, median annual income is highest for respondents in executive or administrative job functions (\$73 000), followed by those in sales or marketing (\$60 240), teaching or training (\$60 000), and R&D (\$59 000). The lowest median income goes to those in design (\$50 000). This group is preceded by those in construction supervision (\$52 000), production or quality control or maintenance (\$52 654), and project study and analysis (\$53 600).

The highest median income was in the Pacific and western states (\$63 400)—up 6.0 percent), and the northeastern states (\$62 000—up 1.1 percent). The lowest median income was in the Great Plains states (\$53 425—up 3.2 percent).

The survey also analyzed income according to 50 metropolitan areas. The highest median income was found in Newark and Jersey City, N.J., and vicinity (\$74 000). It is followed by Nassau and Suffolk counties, N.Y. (\$72 750), and San Francisco and Oakland, Calif., and environs (\$72 000). The lowest median incomes were found in Dayton and Springfield, Ohio, and vicinity

(\$44 575), Memphis and vicinity (\$50 133), and Norfolk, Virginia Beach, and Newport News, Va., and vicinity (\$51 200).

Respondents were also asked, as they had been since 1986, to assess the economic health of their employers. For the first time, engineers reporting that this had worsened since the previous year outnumbered those reporting improved conditions (35.9 to 27 percent). However, asked about their employers' future, 47.4 percent looked forward to an improvement in economic health in 1992, while 15 percent anticipated worse conditions. NSPE pointed out that only 16 percent of last year's survey respondents had forecast the worsening conditions reported by 35.9 percent this year.

The 1992 questionnaire was mailed to 58 615 NSPE members in January 1992. By March 9, 12 792 questionnaires had been returned, a 21.8 percent response rate. After eliminating respondents who were on active duty in the Armed Forces, unemployed or who had been for some part of 1991, full-time students, and retirees, the sample for calculating income was reduced to 11 797. Results of the survey were made available in the late spring.

## Two honored by National Academy

The National Academy of Engineering, Washington, D.C., selected two prominent engineers to receive two of its most prestigious awards. George H. Heilmeyer (F), president and chief executive officer of Bellcore, Livingston, N.J., received the Founders Award for "pioneering contributions to the development of the technology of liquid crystal displays, which have evolved as important components of computers and consumer electronics products, and for visionary leadership in the development of advanced defense technologies."

Ruben F. Mettler (F), retired chairman and chief executive officer of TRW Inc., Cleveland, Ohio, received the Arthur M. Bueche Award "in recognition of outstanding individual achievement in the nation's defense and space technology programs, and for leadership in national defense, industrial trade and competitiveness, higher education, science and technology policy, and public service."

Established in 1965, the Founders Award rewards "outstanding engineering accomplishments by an engineer over a long period of time and of benefit to the people of the United States." The Arthur M. Bueche Award was established in 1982 in honor of Bueche, who served as senior vice president for corporate technology of the General

(Continued on p. 86)



FREE TO IEEE MEMBERS!

This  
Easy-to-use  
Disk<sup>®</sup> Helps  
You *Plan*  
*Your Financial*  
*Future in*  
Less Than an  
Hour!

<sup>®</sup> For IBM PC<sup>®</sup> and compatible systems

NEW

Improved 1992 Edition

- FASTER
- EXPANDED DATA
- PROFESSIONAL ADVICE

Here's how to get your copy of  
the **FREE LifeCycle Review Disk.**

Complete the following information and return this card *today!*

Please print or type.

SIZE OF DISK PREFERRED ☐ 5 1/4" ☐ 3 1/2"

MR ☐ MRS ☐ MS ☐ NAME

ADDRESS

CITY STATE ZIP

DAY TELEPHONE ( ) DATE OF BIRTH / / MO. DAY YR.

IEEE MEMBER NUMBER

65692-9/92

# Is The You Bought ance?



association is  
designed to  
grow with you  
— it even  
moves with  
you when you  
change jobs.

## We're Affordable.

What good would all these benefits be if no one could afford them? That's why we offer members the additional benefit of reasonable rates, negotiated using our group purchasing power. Call 1 800 424-9883 (in Washington, D.C., (202) 457-6820) between 8:30 a.m. and 5:30 p.m. Eastern Time for more information about these insurance plans offered through IEEE:

- Term Life • Disability Income Protection • Comprehensive HealthCare • Excess Major Medical • In-Hospital • High-Limit Accident • Long Term Care • Medicare Supplement • Cancer Expense

## IEEE Insurance

Designed for the way you live today.  
And tomorrow.



# IEEE 488.2

for Your PC



## PC, PC AT, and PS/2 Boards

- 100% IEEE 488.2 compatible
- NAT4882™ IEEE 488.2 Controller chip
  - Optimized GPIB functionality
- Turbo488® performance chip
  - Over 1 Mbytes/sec reads and writes\*
- Includes NI-488.2™ for DOS and Windows software

\* Between boards; adhering to IEEE 488.1 Section 5.2.3.

## Windows and DOS Software

- High-speed IEEE 488.2 routines
- Standard NI-488® functions
- HP-style commands
- Standard Windows DLL entry points
- Interactive control utility
- Controller and Talker/Listener functions

## Call for FREE 1993 Catalog

See us at Autotestcon, booth 605



6504 Bridge Point Parkway  
Austin, TX 78730-5039  
Tel: (512) 794-0100  
(800) 433-3488  
(U.S. and Canada)  
Fax: (512) 794-8411

Branch Offices  
AUSTRALIA 03 879 9422 • BELGIUM 02 757 00 20  
CANADA 519 622 9310 • DENMARK 45 76 73 22  
FINLAND 90 524566 • FRANCE 1 48 65 33 70  
GERMANY 089 714 50 93 • ITALY 02 48301892  
JAPAN 03 3788 1921 • NETHERLANDS 01720 45761  
NORWAY 03 846866 • SPAIN 91 896 0675  
SWEDEN 08 984970 • SWITZERLAND 056 27 00 20  
U.K. 0635 523545

© Copyright 1992 National Instruments Corporation. All rights reserved. Product and company names listed are trademarks or trade names of their respective companies.

Circle No. 22

# The

## Salaries much

Salary increases for engineers were hurt by the recession. The median salary according to the "1992 survey" of members of the Professional Engineering Society, D.C.

The median total compensation for the survey was \$57 000. The figures include income from jobs, bonuses, and other benefits.

The 2.2 percent gain in the index of the consumer price index over the 2.6 percent in 1991 indicates that the engineers had improved their purchasing power.

The highest median salary was in the petroleum industry (\$65 000), followed by chemical (\$62 400), and electrical, environmental (primarily structural) (\$58 000 to \$62 100).

At the low end of the scale, those employed in space (\$56 746), in agriculture (\$55 843), and in manufacturing and least, manual (\$51 185).

As in previous years, the median annual income for those in executive or managerial positions (\$73 000), followed by engineering (\$60 000), and marketing (\$60 000), and lowest median income was in design (\$50 000). Those in construction, production or quality control (\$52 654), and production (\$53 600).

The highest median income was in the scientific and technical services (\$62 000—up 1.1 percent), and the median income was in design (\$53 425—up 3.2 percent).

The survey also indicates that the median income for those in 50 metropolitan areas was \$57 000. Jersey City, N.J., and New York City, N.Y. (\$72 750), and San Francisco, Calif., and the lowest median income was in Springfield, Mass. (\$51 185).

# LifeCycle

REVIEW


**Financial Planning Advice at Your Fingertips in the Privacy of Your Home or Office**

- PROTECT YOUR INCOME
- PLAN YOUR RETIREMENT
- SECURE YOUR FAMILY'S FUTURE

FREE DISK

*Offer to IEEE Members!*

No Postage Necessary if Mailed in The United States



**BUSINESS REPLY MAIL**

First Class Permit No. 10330 Washington DC

Postage will be paid by:

**Administrator  
Group Insurance Program  
for IEEE MEMBERS**  
1255 23rd Street, N.W.  
Washington, D.C. 20077-6916



# Was This The Last Time You Bought Insurance?



It really has been that long. A lot of things have changed since then. The size of your house. Maybe your job. And more than likely, the amount and types of insurance coverage you need. That's why you should have insurance that can easily adapt to the way your life changes — IEEE Group Insurance Program.

## **We Understand You.**

Finding an insurance program that's right for you isn't easy. But as a member of IEEE, you don't have to go through the difficult and time consuming task of looking for the right plans — we've done that work for you. What's more, you can be sure the program is constantly being evaluated to better meet the needs of our members.

## **We're Flexible.**

Updating your insurance doesn't have to be a hassle. With our plans, as your needs change, so can your coverage. Insurance through your

association is designed to grow with you — it even moves with you when you change jobs.

## **We're Affordable.**

What good would all these benefits be if no one could afford them? That's why we offer members the additional benefit of reasonable rates, negotiated using our group purchasing power. Call 1 800 424-9883 (in Washington, D.C., (202) 457-6820) between 8:30 a.m. and 5:30 p.m. Eastern Time for more information about these insurance plans offered through IEEE:

- Term Life • Disability Income Protection • Comprehensive HealthCare • Excess Major Medical • In-Hospital • High-Limit Accident • Long Term Care • Medicare Supplement • Cancer Expense

## **IEEE Insurance**

Designed for the way you live today.  
And tomorrow.



# Forum

## A letter from the Godfather

Beardsley Graham [April, p. 6] has really led with his chin by saying that "Communications satellites were first described by George O. Smith in his novel *Venus Equilateral*, published in 1947."

A little research (may I recommend my perhaps not totally impartial *How the World Was One?*) should have told him that my *Wireless World* article was published in 1945. So obviously I have priority over Smith.

Wrong. *Venus Equilateral* started appearing in *Astounding Stories* in 1942 (how George didn't tangle with Security, I'll never know). I wrote an introduction to the 1965 edition of his book, giving him full credit.

As early as the 1920s, Oberth and other pioneers took space communications for granted—though they often considered using light signals in view of the primitive state of the radio at the time. But to the best of my knowledge, my *Wireless World* paper was the first to point out that satellites were ideal for terrestrial services, that the (well-known) geostationary orbit was the obvious place to put them, and that the power levels required could be easily attained with reasonably sized equipment.

No single person "invented" the communications satellite. It has two fathers—John Pierce (Bell Labs) and Harold Rosen (Hughes). I am content to sign myself—The Godfather.

Arthur C. Clarke  
Colombo, Sri Lanka

## One aspect of diversity

I was pleased to see *IEEE Spectrum* attack the many complex issues posed by diversity ["Diversity in the high-tech workplace," June, p. 26]. But as I read through the section, I grew more and more dismayed to see the concerns—and indeed the very existence—of lesbian and gay engineers almost utterly ignored.

The same day I received the June issue, I also received the latest issue of the lesbian and gay magazine *The Advocate*. It contained an article describing the 10 best places in the corporate world for lesbians and gay men to work. Fully half of those companies are in the business of electrotechnology: AT&T, Digital Equipment, Lotus Development, US West Communications, and Xerox. Together these companies employ more than half a million people.

Each of these five companies explicitly bans anti-gay bias in hiring and promotion. Most have employee health groups for les-

bian and gay workers, and some offer health or other benefits to employees' same-sex domestic partners. Last year, Lotus Development Corp. became the first publicly held company in the United States to offer full spousal benefits to gay men and lesbians.

Surely these examples of humane and progressive employee policies at technology companies should have been included.

John Voelcker  
New York City

## A sharper focus on fuzzy logic

To those who may still be doubtful about the idea of fuzzy logic ["Fuzzy logic flowers in Japan," July, pp. 32–35], it might be pointed out that the idea in primitive form is very old, namely, in on-off control.

A typical example is an on-off thermostat, or in more advanced form, an on-off thermostat with two heaters, which is closer to the typical fuzzy logic situation. If the two heaters are of 100 and 200 W, then when the signal from the temperature sensor is very low, both would be switched on. When the signal is low, the 200-W heater would be switched on; when it is a little low, the 100-W heater would be on; and of course, when the signal is high, they would both be off.

Joseph M. Diamond  
Brooklyn, N.Y.

## On the CASE

The following statements in "CASE's missing elements" [June, p. 39] are not factual:

- "DOD-STD-2167A... requires the creation of more than 50 different document types." (The maximum is 20, and in most defense system applications developed under this standard, the number is much less).
- "Indeed, more than 50 percent of the entire cost of U.S. military software is devoted to paperwork, while pure coding accounts for less than 20 percent."

While one generally agrees with the six chronic problems of software engineering that make CASE [computer-aided software engineering] especially challenging, I would like to add one more to the list: the misapplication of software engineering standards and the promotion of related folklore and horror stories. Jones's article serves well to demonstrate this point—and with due respect to defense software practitioners, needs to be corrected.

There are an estimated 300 000 software and system engineers in the United States developing defense software and some of the

most complex real-time and unprecedented systems in the world. Their work requires a disciplined development approach and considerable coordination of the total effort.

The end product of this intellectual activity is information—electronic or hard copy. Even though the end product may appear to the casual observer to be "paperwork," most of this effort is or should be expended on engineering the product.

During development of defense systems, the DOD-STD-2167A is the contractually imposed development standard. Since 1985 the DOD-STD-2167 has defined uniform requirements for this market segment and facilitated transition from *ad hoc*, project-specific processes to standardized and mature software engineering processes. This market segment, due to its complexity and challenges, has been the primary target for the CASE industry and the initial sponsor for many of its products.

The DOD-STD-2167A is a highly flexible third-generation standard that has evolved since the 1970s. It defines the requirements for a mature software process. It represents the joint product of DOD [Department of Defense] and industry, including over 500 contributors, 18 000 comments, the resolution of 55 major issues, and 12 drafts. The standard continues to evolve as a function of changes in industrial practices and technology. Current efforts under way include ■ *DOD-STD-2167A Rationale and Application Handbook* and the development of DOD-STD-2167 Revision B.

The DOD-STD-2167A development concept is based on a unified, streamlined, and tailored Software Development Library (SDL). The SDL information content, in contrast to the horror story in the article, contains ■ maximum of 20 and not 50 data items. The SDL information is structured into five areas: management and planning, requirements specifications, design, test, and user support. Depending on systems complexity and its life cycle phase, as few as seven data items may be required. Contrary to widely held misconceptions, the DOD-STD-2167A specifically encourages CASE and electronic delivery of information. There are examples of well-tailored DOD contracts where only electronic access to SDL information is required—with no hard copy delivery or paperwork.

Any standard can be misapplied. This is particularly true of DOD-STD-2167A because of its flexibility and the need to tailor its requirements across the wide spectrum of DOD system applications. The lack of a concurrently issued *Rationale and Applica-*

(Continued on p. 85)



# Introducing Power Windows For Project Managers.

*The #1 Rated Project Manager Now Available For Windows.*

Power changes people.

Especially project managers.

CA-SuperProject

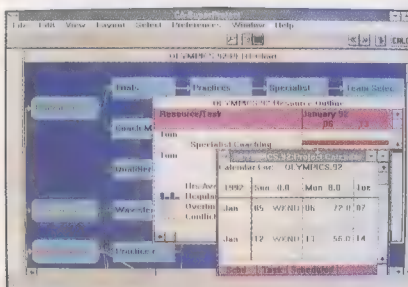


They're working smarter and faster with new CA-SuperProject® For Windows.

It's the world's most advanced, efficient and reliable project management software

— and now it's incredibly easy to use.

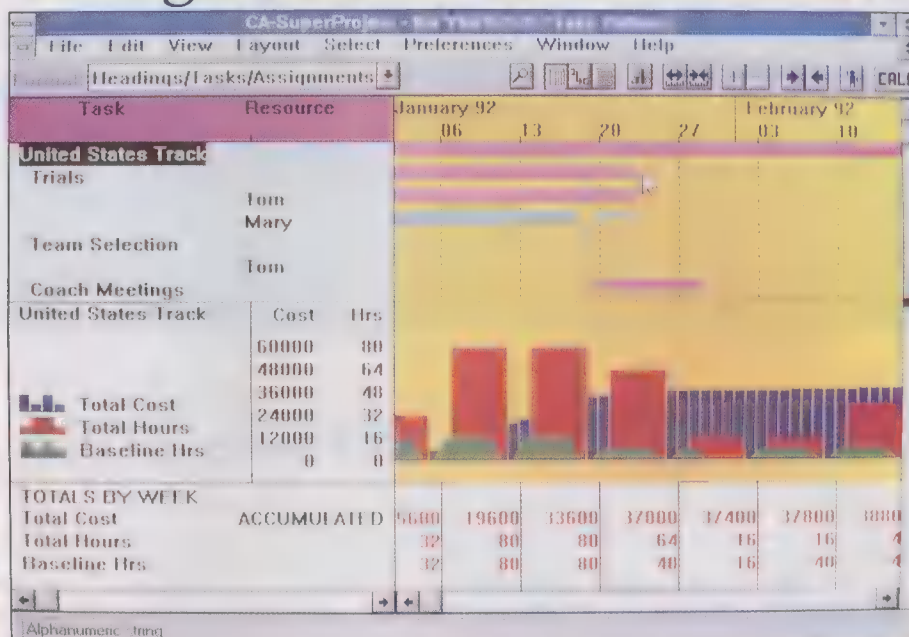
Total power is yours with just a few mouse clicks. Create and edit projects. Specify resources, task types and durations. Define integrated sub-projects. Build top-down hierarchies and task-



Report the status of your project with detailed Gantt, PERT, WBS and Cost/Resource charts.

dependency relations. Link multiple projects together for cross-project leveling. Perform extensive "what-if" analysis, revising schedules as projects progress.

You can bet your career on its



Manage with power, using tools for sophisticated planning, comprehensive resource management, tracking and controlling.

advanced and efficient scheduling algorithms.

A recent study of the five leading project managers proved it. Each was assigned the same



ing days—leaving Microsoft Project, Timeline 4.0, Project Workbench and Project Scheduler in the dust.

There's also a wide array of state-of-the-art graphics and detailed reporting tools to help bring your projects to life.

For your free Demo Disk, call 1-800-CALL CAI. Call today.

And find out what our power windows can do for you.

**COMPUTER ASSOCIATES**  
Software superior by design.

**CA-SuperProject® For Windows**

© Computer Associates International, Inc., One Computer Associates Plaza, Islandia, NY 11788-7000. All product names referenced herein are trademarks of their respective companies.

Show multiple views of the same project or different projects simultaneously.

project, but the finish dates varied by as much as five months. CA-SuperProject For Windows finished first in 214 work-






Image: Earth's atmosphere on Oct. 1, 1987, correlating temperature (bluish haze) with wind velocity, direction and horizontal speed (colored arrows) at 1000-100 mb. Data courtesy of NASA/GSFC/NSDC

# Your workstation will think it's dreaming.

Put the visualization environment of the IBM visualization supercomputer on your SGI, Sun, HP or IBM RISC System/6000 workstation.

IBM Visualization Data Explorer delivers all the problem-solving functionality of IBM's POWER Visualization System, tailored to workstation capacity.

What distinguishes Data Explorer from other visualization packages is that it lets you focus on visualization — without getting bogged down in technology. It allows you to freely combine visual techniques within the same image.

Data from various sources, such as simulations and observations, can be quickly viewed together. And with Data Explorer's ability to directly handle time-series, producing exciting animations doesn't have to be a nightmare.

Data Explorer also has a powerful, extensible tool kit that provides fast access to a broad range of sophisticated functions...both visual and computational.

Whether you're a visualization novice or expert, you should know more about IBM Visualization Data Explorer. Call (800) 388-9820. Or write: IBM Scientific Visualization Systems, T. J. Watson Research Center, P.O. Box 704, Yorktown Heights, NY 10598.



# Spectral lines

SEPTEMBER 1992 Volume 29 Number 4

## It's the law, but is it ethical?

**W**e are told, and it seems true, that the United States leads all nations in lawyers per square centimeter.

Corporations could not survive without them, as they must protect proprietary technology and fend off lawsuits by competitors, disaffected stockholders, unhappy customers, vendors and employees, and probing government agencies.

Two specialties of the law—patent law and liability law—were abetted by developments in high technology.

Some of us, when we were young engineers, believed that patent law, unlike many other areas of civil law, was based on fact and logic, not on emotion or half-truths. After all, patent law concerned science and engineering, did it not? And mathematical and physical laws. We believed the same of liability law.

However, the reading of history soon taught us otherwise. And corporate lawyers, who are paid a great deal of money to restrict our options, instructed us to the contrary. Above all, they said, survival of the corporation comes first. But, we asked, does that mean, if we are called to testify on a subject of concern to our company, we should shade our testimony, tell half truths, mislead or even lie? "The corporation comes first," they repeated.

If we are to believe the historians of technology who spend most of their time researching these matters, even eminent, respected pioneers in radio and electronics were convinced by corporate counsel to express unprovable opinion as fact in order to suppress competition, or to gain an economic advantage for their own companies.

As just one example, when the Radio Corporation of America (RCA) became concerned about the threat of frequency modulation to the entrenched and profitable amplitude modulation radio broadcasting, it lobbied before the Federal Communications Commission against allocation of favorable frequencies to the upstart FM. Charles Jolliffe, head of RCA's engineering department and a former commission chief engineer, testified in 1936 concerning frequency allocations and failed completely to mention FM, an oversight interpreted by Edwin Armstrong, the developer of FM, as outright conspiracy.

RCA patent attorneys also put the work of one of their own engineers, Murray Cros-

by, into interference with Armstrong's. Tom Lewis, in his recently published *Empire of the Air*, noted that Crosby's own notebooks revealed an FM circuit diagram that Armstrong himself had given to Crosby.

If all this is true—and not much evidence has been forthcoming to dispute it—what does this imply about the values of engineers today? Is the situation better or worse?

Are the economic pressures and career constraints so great for employed engineers that the lawyers' admonition that the corporation's interests are supreme prevails over good science and good engineering? Are engineers so intimidated that good science/engineering is subverted to falsely depict a competitor's product as inferior or not worthy of a patent?

Perhaps it should not be surprising that only one engineer among many privy to the dangerous design and questionable test results of the Challenger booster rockets would dispute the decision to launch NASA's ill-fated space shuttle.

In today's culture, in the United States and apparently in most of the rest of the world, the prevailing ethic seems to condone deception. "If you can get away with it, it's okay" applies to professional sports (a foul missed by the referee is okay); much of the law (a defense lawyer is expected to depict a victim as unworthy of winning a verdict); financial transactions (a seller discloses only what the law requires, and charges for products or services what the traffic will bear, especially if it's reimbursed by medical insurance); and ownership of property (it's the owner's fault if the car was stolen, especially if the key was in the ignition).

Unfortunately, many engineers seem to assume the coloration of their environments, and easily adapt to the corporate lawyers' culture. The umbrella of "fiduciary responsibility" may be tenuously extended over them by the corporation, intimidating them into actions foreign to their inclination and training.

Of course, the lawyers are not totally to blame. Engineers helped create the field of patent law by inventing complex, arcane devices and equipment that required definition, documentation, and protection so the inventors (or their companies!) could profit. But, as with many other aspects of technology, engineers were unable to control the practice of patent law. It evolved for better or worse.

If you deem the foregoing to be mild criticism and cautionary concern for engineers, the situation may be even worse in the area of liability law. There the lawyers themselves express skepticism about the legitimacy of expert technical witnesses.


In his recent book, *Galileo's Revenge*, author Peter Huber (himself a former law clerk to Supreme Court Justice Sandra Day O'Connor) comments on expert witnesses in liability cases. Attorneys never seek objective scientists or engineers, he writes. Rather, they look for a strong bias.

Huber writes: "Why would you, the diligent lawyer, settle for a scientist who will say that PCBs may in some circumstances affect health, though how and at what concentrations is most unclear, if you can find one who will swear that they are one of the most lethal substances known to man, that they subvert the immune system, and they undoubtedly were to blame for this plaintiff's migraine headaches? Why settle for one who will say that 60-cycle [hertz] electromagnetic fields probably don't injure human health, though one must concede certain small pieces of disquieting evidence to the contrary, if you can find one who will... absolutely, positively promise that the fields will do no harm, no how? The middle of the road, in law even more so than in politics, belongs to the yellow stripes and the dead armadillos. So, as you labor to assemble your case, the strength of the scientific support for an expert's position is quite secondary. It is the strength of the expert's support for *your* position that comes first."

Huber also suggests that engineers exposed to the courtroom environment, as expert witnesses, may get caught up in the excitement of the legal theater, become advocates of the lawyers who hired them, and find the taste of winning overriding their respect for the truth.

Ethics are often not clear-cut, but falsification for economic gain, however rationalized, is not in the gray area.

Donald Christiansen





**SPECIAL REPORT: SUPERCOMPUTERS**

# **TERAFLOPS GALORE**





S

upercomputing, a field of endeavor perhaps 20 years old, is entering ■ new era.

Born in the wilds of the northern mid-west United States and nurtured by esoteric Cold War pursuits like nuclear weapons design and code-breaking, the conventional supercomputer now seems poised for an indefinite but inexorable decline. On the verge of a gradual takeover, industry analysts believe, is the massively parallel processor (MPP), which is already winning a place for itself ■ the high-performance architecture of choice.

At the same time, supercomputers are growing beyond their roots in the military, intelligence, and scientific communities, and into mainstream business, industrial, and design applications. Already, supercomputers have been used to design automobiles that will better protect passengers in crashes, internal-combustion engines that burn fuel more efficiently, pharmaceuticals that are safer and more effective, and integrated circuits of unprecedented complexity. There are thousands of commercial applications now in use, and many more envisioned for much more powerful machines that are still years away.

The converging trends in supercomputing, which foreshadow ■ transformation in how supercomputers are built and how they are used, will profoundly alter the way products and materials are designed, the way research and experiments are conducted, and even the lifestyles of people in developed countries. Most economists are now convinced that the financial success of nations will depend on their ability to innovate, design, and manufacture—three realms in which the role of supercomputers is important and becoming dominant.

Quietly abetting this revolution is the MPP, ■ computer architecture based on the interconnection of hundreds or thousands of microprocessor-based nodes (as opposed to the four, eight, or sixteen expensive, exotic processors in ■ conventional supercomputer). With ■ degree of uniformity that is rare in computing, the market's manufacturers have concluded that massive parallelism is their best hope of providing what users want: processing rates of ■ trillion floating-point operations per second, sustained on actual applications. This is 100 times the processing power of today's best machines.

It is ■ tall order. Only in the last few months could a handful of exceptionally lucky users even hope to run their applications at about 10 billion floating-point operations per second (10 gigaflops). There are, after all, fewer than 500 conventional supercomputers in the world and perhaps a like number of powerful MPPs, and only in the last year or so has the tiniest fraction of machines in either category begun approaching or exceeding the 10-gigaflops rate.

A teraflops, at least, would let researchers simulate, with reasonably fine detail, ■ variety of phenomena they have never been able to model before. In this category would be the aerodynamics of ■ entire aircraft, the global climate over a period of decades, and the formula and atomic structure of ■ advanced material with certain desired properties. Incredibly, at least a half-dozen manufacturers are confident they will be building MPPs that can sustain a teraflops on such problems before the decade is out.

Glenn Zorpette Issue Editor

*The electron density inside a water molecule was calculated by IBM Corp.'s power visualization system. The colored contours show the probability of finding electrons; arrows show the magnitude and direction of the gradient of electron density.*

IBM Corp.

Such ambitions have stirred up considerable technical ferment and ■ diversity of architectures [opening article, p. 28]. But perhaps more than for any other computer type, software is critical to high performance. From the outset, software has been viewed as the Achilles' heel of MPP technology, though Richard Comerford notes that long strides have been made lately [p. 34].

Given the importance of software, shouldn't MPP development be paced so that software and hardware grow and evolve together? George Cybenko and David J. Kuck think so; unfortunately, they note, MPP vendors, the U.S. government, and even buyers of the machines seem to have other ideas [p. 39].

National security is no longer their *raison d'être*, but supercomputers are no less national resources than they were in the Cold War—only today, high-performance computing is more likely to be used to compete economically rather than posture militarily. This situation helps explain why the U.S. and Japanese governments treat their supercomputer industries as they do—and why both U.S. vice presidential candidates made it a point to visit Eagan, Minn.-based Cray Research Inc., the world's premier supercomputer maker, while campaigning last summer.

Japanese supercomputer makers are unsurpassed in the construction of very fast semiconductors and single processors, and will soon plunge into the MPP market, according to David K. Kahaner and Ulrich Wattenberg, two Tokyo-based analysts who study Japanese supercomputing [p. 42]. Recent achievements in that country include ■ record-setting benchmark of 20 billion floating-point operations per second (gigaflops) by NEC Corp.'s SX-3; the fabrication of an air-cooled gallium arsenide chip, also by NEC, with a peak performance of 0.2 gigaflops; and Hitachi Ltd.'s new S-3800 supercomputer, whose claimed peak of 32 gigaflops is the world's highest.

Surprisingly, however, Japan has not managed to network its supercomputers to the extent the United States has. And as supercomputers become more powerful and insinuate themselves into mainstream uses, the accessibility provided by networks becomes indispensable [p. 48].

Supercomputers are already being used to design everything from turbines to toilet bowls. A number of vital scientific and technical "Grand Challenges" awaits more powerful machines [p. 56]; sophisticated visualization techniques will help users see how giant programs execute and make sense of the huge sets of data they produce. On pp. 61–65, Matthew Arrott and Sara Latta explain the field's basics and speculate about its future.

Supercomputing is by definition computing's cutting edge, the proving ground for technologies that will eventually trickle down to less powerful and less glamorous kinds of machines. Exotic forms of silicon and gallium arsenide semiconductors are being developed for the next generations of supercomputers. But

researchers believe that electronics may eventually have to give way to optics or some other technology if processing rates are to rise well beyond a trillion floating-point operations per second. At any rate, it takes only a little imagination to see the concluding section, "Future technologies" [pp. 66–75], ■ ■ rough diagram of the 21st century computer.

There are rare moments in technology when the next 5, 10, or 15 years begin to crystallize; this is one of those times. Of course, as with any undertaking of this magnitude, there will be many unexpected twists and turns: innovative architectures, software breakthroughs, and as-yet unimagined applications. From such surprises will come much of the excitement ■ high-performance computing reaches new heights. ♦

## REINVENTING THE MACHINE

- 28 The power of parallelism
- 34 Software on the brink
- 39 Revolution or evolution

## NATIONAL EFFORTS

- 42 Japan: ■ competitive assessment
- 48 Balancing resources

## EMERGING APPLICATIONS

- 56 Modeling reality
- 61 Perspectives on visualization


## FUTURE TECHNOLOGIES

- 66 Superchips for supercomputing
- 69 Interconnections and packaging
- 72 Beyond today's supercomputers
- 76 To probe further



# The power of parallelism

*Extraordinary technical ferment is the rule as supercomputer makers jockey to be first to deliver extremely fast machines*



Like ■ young adult venturing out into the working world for the first time, the world's supercomputer industry is headed for an unknown but exciting new phase in its life. As major supercomputer applications shift increasingly from the sheltered environments of Government and academia

to the less forgiving arenas of business and finance, the machines themselves are moving toward processing rates above tens of billions of operations per second, up from millions or billions per second. This kind of performance will have far-reaching effects on many of the most critical enterprises undertaken by advanced nations, such as monitoring and forecasting climate globally and designing better aircraft, vehicles, engines, materials, semiconductors, drugs, and medical treatments.

But as with any transition of such a magnitude, there will be growing pains, and they will not be trivial. Both triumphs and travails will be inextricably linked to the computing technique that brings them about: massive parallelism, in which machines are based on hundreds or thousands of processors not much different from those at the heart of high-performance workstations.

The technique has been around for more than a decade, but only recently has the industry reached a consensus that it is the only way to build machines capable of the exalted processing rates that many users are demanding—rates as high as a trillion floating-point operations per second (teraflops). At present, however, there is not much difference between the performance that the most powerful massively parallel processors (MPPs) and the most powerful conventional supercomputers can sustain on actual problems. Both have managed several billion floating-point operations per second.

Glenn Zorpette Issue Editor

The crucial difference is that some MPP architectures are already said to be scalable to ■ teraflops, although it will be several years at least before such processing rates can be had for ■ realistic price—US \$30–\$40 million, say. Hopefully, by that time, more advanced software will be available to help users milk the most out of the machines, because many agree that, despite recent improvements, programming them to run efficiently remains a challenge.

“Death, taxes, and parallelism: nobody’s in favor of any of them, but they’re inevitable facts of life,” said Paul Schneck, director of the Supercomputing Research Center (SRC) in Bowie, Md. One chief executive of a company that makes MPPs recalled an executive of ■ Fortune 500 company telling him that the machines should be called “massively fraudulent” because “only the DOD [Department of Defense] can afford to program those things.”

**NEEDED: STANDARDS.** Greatly complicating the move toward MPPs is the diversity of approaches and consequent lack of standards. “The industry has to settle down—we have a lot of companies coming and going relatively quickly,” said Larry Rapagnani, associate vice president for computing and information technology at the University of Arizona in Tucson. “People are fearful of expending ■ great deal of time writing software without knowing if the vendor will be around for ■ while.”

Unprecedented  
computational  
power will transform  
critical industries  
and technologies

In the United States, at least eight companies offer high-performance parallel computers, and at least a half dozen more—including heavyweights like IBM, Cray Research, and Convex Computer—are planning to introduce MPPs within ■ year or two. However, a number of ventures have already failed, even though MPPs as an industry segment are only about six years old. The failed efforts include MPP ventures at Alliant Computer Systems, Bolt Beranek & New-

man, and Ametek in the United States; Suprenum in Europe; and Myrias in Canada. In July, Active Memory Technology Inc., an Irvine, Calif.-based MPP company, filed for protection under chapter 11 of the bankruptcy code.

According to the Smaby Group Inc., the leading supercomputer market researcher, sales in the MPP market reached about \$270 million in 1991, vs. about \$1.2 billion for conventional (vector) supercomputers. But according to the Minneapolis, Minn., firm, the MPP market is expected to experience compound annual growth of nearly 35 percent until 1996—roughly four times that predicted for the vector market. At that rate, the MPP market will more than double by 1994 to sales of over \$600 million.

Still, it is doubtful that more than a few of the 15 or so current and near-future MPP makers in North America will be successful—especially as there are only about ■ half-dozen makers of vector supercomputers in the world, and not all of them are turning ■ profit on their supercomputer lines.

**TECHNICAL UPEHAVAL.** Competition in the budding market has engendered considerable technical ferment. For the competitors, the challenge is scalability—designing an extremely high-performance architecture capable of accommodating thousands of processors *without* suffering from bottlenecks caused by insufficient memory or communication between processors.

Until recently, discussions about massively parallel architectures often boiled down to debates over the relative merits of single-instruction, multiple-data (SIMD) and multiple-instruction, multiple-data (MIMD) machines. With SIMD, all the processors execute the same instructions in lockstep, but on different data. Since the processors are always synchronized and each handles a different piece of data on each cycle, thorny questions about the status of ■ piece of data and which processors may modify it are largely avoided. So the machines can often be built more simply and inexpensively than comparable MIMD designs.

The usual criticism of the SIMD model, though, is that it works best on a certain class of problems—in image processing, for example, where the same operation is performed on every byte representing every pixel in an image with ■ million or more of them. Perhaps because of this fact, most



conventional mainframes and supercomputers.

Distributed computing is the loosest form of parallelism, in which parts of a problem are divided among various nodes and there is relatively little communication once processing begins. This is the way a cluster of workstations or minicomputers would attack a single problem, for example.

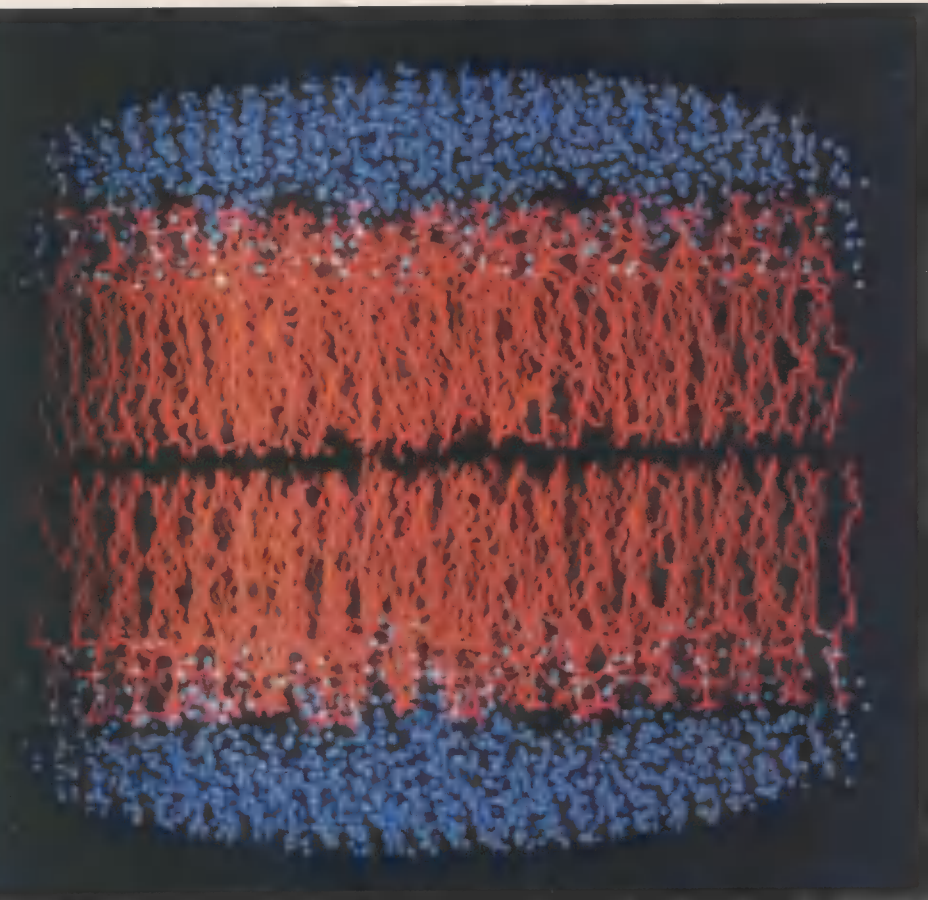
With message-passing, the many processors in an MPP cooperate in solving a large problem by sending messages of data and code to each other. The communication method is like packet-switching: every processor in the system has an address, and routing information in a message guides it through the network, sometimes through intervening processors, to its destination. There may or may not be communication between processors to ensure that the message was received. Generally, the routing of messages is explicitly worked out by the programmer, possibly with the assistance of various software tools, as the program is written. Message-passing is the most popular programming model for MPPs.

Data parallelism has become a nebulous term for programming models that are neither message-passing nor distributed. The term once referred to the way SIMD machines are programmed, but several MIMD machines of late have added hardware and software to enable users to program in SIMD style. Programs are not necessarily executed in lockstep, however; synchronization may be broken wherever convenient to let processors run independently.

**TWO MODELS IN ONE.** Thinking Machines Corp., Cambridge, Mass., began as a maker of SIMD computers, but its current model, the CM-5, lets users program in message-passing or data-parallel modes, or a combination of the two. Processors are connected in a fat tree, so named because of an arboreal analogy. In this topology, the processors are like leaves on a tree, and the communication links like the branches—they are “fatter” (have more bandwidth) when they are farther from the processors, thus connecting more of the processors together.

At least, this is the ideal. In practice—in the CM-5, for example—the fat tree is really thin, on the assumption that most communication will be local. So the branches need to be thickest when connecting nearby processors, not distant ones.

Key to the CM-5's dual programming modes are its two separate communications networks linking all nodes in the machine. (In an MPP, a processor and all the associated hardware that lets the processor communicate with its memory and other processors is often referred to as a node.) One of the CM-5's networks is used to transfer data between nodes, which is critical for message-passing. The other network is used for operations involving all the nodes, such as synchronizing or coordinating them or broadcasting data or instructions to all of them;



Researchers at the Beckman Institute for Advanced Science and Technology at the University of Illinois, Urbana-Champaign, are studying the dynamics of cell membranes using various massively parallel computers, including a home-made system based on the transputer microprocessor and a processor board from Parsytec GmbH. The image depicts a two-layer membrane [red] in a nonspecific biological cell with water [blue] on both sides of the membrane. Each pair of the red strings represents a single phospholipid molecule in the membrane, while the blue dots represent water molecules. In all, 23 978 atoms are depicted.

MPP makers striving to build general-purpose computers have settled on MIMD architectures. A few companies are still pursuing the SIMD approach, however, including MasPar Computer Corp. in Sunnyvale, Calif., and Wavetracer Inc. in Acton, Mass.

In a MIMD machine, as well as in an ordinary mainframe or supercomputer, for that matter, the processors may be pursuing a common goal or solving a single problem, or they may be solving many problems at the same time. In a MIMD machine, however, each processor has its own memory and can be executing a different program. In the more conventional machine (also known as a symmetric multiprocessor), all the processors access a large common memory. As of now, there are technical limitations to how many processors can be built into the symmetric kind of system, so all the current MPPs are distributed-memory machines.

Nonetheless, the relative ease of programming symmetric multiprocessors, with their shared memories, has led several companies—notably Kendall Square Research (KSR) Corp. in Waltham, Mass.—to incorporate some aspects of this approach into their distributed-memory architectures.

Indeed, some industry experts see this as the wave of the future in massively parallel architectures (more on that later).

“A second phase is coming that everyone is going to have to react to to survive,” said Stephen Nelson, vice president of technology and director of Cray Research Inc.'s MPP project. Eagan, Minn.-based Cray is one of several companies developing an MPP which, like KSR's, grafts a shared-memory model onto a distributed-memory architecture.

**GRIDS AND TREES.** MIMD MPPs may also be classified according to topology and programming model. Popular topologies, or ways in which processors are interconnected, include the so-called fat tree, in which processors are grouped into clusters or clusters of clusters; and the “mesh,” in which processors are arranged in a (typically) two-dimensional grid. The most common programming models, which are determined primarily by topology and the way processors communicate, are: distributed computing, message-passing, and data-parallelism. A fourth, emerging programming model gives users some of the simplicity of the shared-memory architecture of



thus it is important for data-parallel programming. Both networks ■■ scalable, their bandwidth increasing as the number of nodes increases.

Each node in ■ CM-5 has 8, 16, or 32M bytes of memory and is connected to the node's processor by ■ 64-bit-wide bus. Each node also has up to four vector units—a unique option among MPPs—although as of July, Thinking Machines had not yet delivered any of the vector processors to customers. According to the company, the maximum complement of vector units boosts the peak double-precision performance of ■ single node from 5 to 128 megaflops, and peak memory bandwidth within the node from 128 to 512 megabytes per second.

Each vector unit is controlled by the node processor, ■ 33-MHz Sun Sparc microprocessor, which fetches instructions for itself and the vector unit. At least one user of the CM-5 wondered aloud, during ■■ interview, how the machine's 5-megabyte-per-second node-to-node communications links could possibly keep up with nodes racing along at 128 megaflops.

**100+ GIGAFLOPS BY JUNE.** The other titan in the MPP market is Intel Corp.'s Supercomputer Systems Division, Beaverton, Ore. Although the division does not release annual sales figures, it may be the largest MPP maker by this measure [see table, p. 32]. Together, Intel and Thinking Machines supply about two-thirds of the MPP market.

The two organizations have more in common than relative commercial success. Like the CM-5, Intel's machine, called Paragon, is a MIMD machine that supports both the message-passing and data-parallel programming models. The chief difference between the machines is in topology: the Paragon's processors are connected in ■ two-dimensional mesh (an ordinary flat array) rather than ■ fat tree.

Paragon is a direct descendant of the Touchstone Delta, installed at the California Institute of Technology (Caltech) in Pasadena in May 1991. Owned and operated by a consortium of Government and academic laboratories, the Delta is by almost any standard one of the world's most powerful computers. In November 1991, the 528-processor machine achieved a record 13.9 gigaflops on a highly parallel Linpack benchmark, which involves solution of ■ large system of linear equations. The record stood until last spring, when the NEC SX-3, a vector supercomputer, achieved 20 gigaflops on the same benchmark, according to Jack Dongarra of the University of Tennessee in Knoxville and Oak Ridge National Laboratory, Oak Ridge, Tenn. For comparison, Cray Research's Y-MP C90 achieved 13.7 gigaflops on the highly parallel Linpack, Dongarra said.

Last July, Intel announced that it had received a contract from Oak Ridge National Laboratory for ■ 512-node Paragon worth approximately \$40 million. The machine, with

■ projected peak processing rate in excess of 100 gigaflops, is to be completed in the late summer of 1993.

The Paragon's computational nodes will each have either two or five Intel i860 XP microprocessors and 16–64M bytes of memory (the 512-node Oak Ridge machine was scheduled to be installed this month with two microprocessors per node, and then upgraded by next summer to five microprocessors per node). In every node, one microprocessor will prepare and send messages to other nodes, while the rest will act as computational engines. All nodes will also have ■ router, which receives messages, checks their addresses, and then either accepts them or passes them on to any of the four nearest nodes. The current Paragon design supports from 66 to 1000 computational nodes, according to Peter Wolochow, the division's market segment manager for Grand Challenge computing.

**COMPUTING IN 13 DIMENSIONS.** NCube, one of the oldest MPP makers along with Intel and Thinking Machines, has made no drastic changes to its basic architecture but is nonetheless finding itself farther and farther from everyone else in the MPP pack. Its current-generation machine, the nCube 2S, is based on custom-designed processor nodes, linked in a hypercube topology. Both features have been tried and abandoned by other vendors, but in an interview, nCube's president, Michael Meirer, argued that the combination of the two was a winning one.

Developed at Caltech in the early 1980s, the hypercube was the topology of choice for the first crop of commercial MPPs introduced around 1985. With this topology, the number of nodes is always a power of two, and all have dedicated communications to their nearest neighbors.

## Massively parallel machines capable of 100 billion floating-point operations per second are being built

The design is ■ elegant one, and inherently scalable: the number of nearest neighbors, and thus communication channels, goes up as the number of nodes goes up. Nonetheless, the architecture has its drawbacks: if the machine is to be expanded and the hypercube topology retained, the number of nodes must be doubled—an expensive proposition if the machine already has 1024 nodes. Also, internode communications pose a dilemma: if the paths connecting nearest neighbors are all bit-serial, communications are often inadequate.

On the other hand, if the communication

paths are ■ byte wide or wider, a relatively huge number of semiconductor pins are needed to implement a hypercube with hundreds or thousands of nodes, and such a system can be difficult to make highly reliable. Partly because of these factors, Intel and Thinking Machines, whose first offerings were hypercubes, have retreated from the topology.

"We think there's one very good reason for them to back away," said Meirer. "Unless you have all the hardware drivers and communications channels on chip, you cannot implement them. You end up with too much hardware; the reliability and sheer size become unmanageable."

The node of ■ nCube 2S, the current model, has two elements: ■ 0.8- $\mu$ m, half-million-transistor, custom CMOS chip; and up to 64M bytes of memory. Integrated onto the custom chip is a floating-point computational unit, a memory management interface, a message-routing unit, and 14 pairs of direct-memory-access (DMA) channels. Thirteen of the pairs of DMA channels may be used to connect the node to its nearest neighbors (the fourteenth is for system input and output), so that the maximum number of nodes is  $2^{13}$ , or 8192 (the largest configuration that nCube has sold, however, is 1024).

**CUSTOM CHIPS OR OFF THE SHELF?** With its highly integrated, custom ICs, nCube can squeeze 64 nodes onto a single circuit board, and ■ 1024-node machine into a chassis with a footprint of just 1.58 m<sup>2</sup>. But a custom processor has drawbacks, too. The competition among makers of reduced-instruction-set computing (RISC) microprocessors is fierce, with ever more powerful—and less expensive—chips being introduced almost every other month. An MPP line designed

around one off-the-shelf RISC chip can generally be easily and quickly modified to use a more powerful one ■ as soon as it becomes available.

"If you're not taking advantage of these high-volume technologies, you're missing the boat," said James Bailey, marketing director at Thinking Machines.


Henry Burkhardt, president of Kendall Square Research, couldn't disagree more. His company's KSR1 is based on ■ custom processor comprising four CMOS chips, which are fabricated by

the Japanese electronics giant Sharp Corp., Osaka. Sharp uses high-volume, low-cost manufacturing techniques developed for the consumer electronics industry. Burkhardt insists that the cost is "lots lower than if it were a merchant-supplied chip."


"The point is, this is ■ chip designed to do supercomputer work, not for the desktop," he said.

**VECTOR VS. MPP.** While such debates rage on among MPP designers, conversations with MPP users make it clear that they have ■ somewhat different set of concerns centered on efficiency, performance, and ■ of use.








Hypercube



Two-dimensional mesh



Fat tree

Among the factors that distinguish competing massively parallel processors (MPPs) from each other is topology, or the way in which processing nodes [shown here in red] are interconnected via communications links [black]. When the first MPPs were released in the mid-1980s, the versatile but technologically demanding hypercube was the architecture of choice, but lately meshes and fat trees  to be gaining in popularity. This diagram shows the ideal fat tree, in which communications links have more bandwidth (are "fatter") when they interconnect more nodes. But in practice, they  usually thinner, on the theory that nearby nodes need to communicate more than faraway ones.

For example, although most MPPs have astoundingly high peak execution rates, in the tens of gigaflops or higher, their performance on applications and realistic benchmarks is often no better—and sometimes worse—than vector supercomputers with much lower peak rates.



In a series of comparisons, David H. Bailey of the National Aeronautics and Space Administration's (NASA's) Ames Research Center at Moffett Field, Calif., found that the ratio of sustained performance to peak performance is typically 1-5 percent for Thinking Machines' CM-2 and Intel's iPSC/860, and 30-60 percent for the Cray Y-MP. Bailey also found that "when sustained performance rates are normalized by system prices... the highly parallel systems are approximately on a par or slightly less cost-effective than the Cray systems."


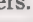
Such criticism is vehemently rejected by Richard Clayton, vice president of Thinking Machines. "It's a crock," he declared impatiently in an interview. According to Clayton, the sustained-to-peak ratio is "a totally worthless metric."

"Having a little extra peak around is a good engineering choice for building a cost-effective supercomputer," Clayton insisted. "Fully utilizing peak processing is important if that's the costly resource," which it is not, he noted, thanks to mass production of RISC microprocessors.

Wendy Vittori, director of marketing and strategic planning for Intel's supercomputer division, also criticized the NASA Ames study, saying it "did not focus on optimizing the performance of parallel computers." As evidence of the efficiency attainable with MPPs, she cited the fact that in 1990 a team

using an Intel iPSC/860 won the annual Gordon Bell prize for exceptionally low price/performance ratios in high-performance computing.

The team, from two U.S. government laboratories and an Italian university, coaxed 2.5 gigaflops out of  128-processor machine—35 percent of its peak rate. The high rate was achieved by redesigning the algorithm, which was used to study the electronic structure of materials such  superconductors, and by recoding some routines in assembly language.


**MASSIVELY PROBLEMATIC?** This achievement, however, also illustrates  lingering weak point: ease of use. The Gordon Bell prize was won not only by the Intel machine but by talented programmers who redesigned an algorithm and wrote programs in assembly language—skills not possessed by all users of scientific and technical computers. To be sure, these skills may not be necessary merely to use  MPP, but they seem to be essential to getting the most out of it.



Boeing Co. in Seattle, Wash., has four Intel MPPs—no commercial organization has more—and plans to replace one of the four with a Paragon. How does programming an MPP compare with programming a conventional supercomputer? "The message-passing paradigm is not terribly difficult to use, but it's hard to get a lot of performance out of," said John Lewis, associate technical fellow at Boeing Computer Services, Bellevue, Wash.


The upshot: "It's considerably harder to use than a vector machine." Lewis's experiences with Thinking Machines' CM-1 and CM-2 (predecessors of the CM-5) con-

vinced him that it is "a lot easier to use [the data-parallel paradigm], but even harder, if not impossible, to get high performance out of it."

One way of getting the most out of an MPP in message-passing mode is overlapping communications, Lewis noted. With messages being sent and received simultaneously among nodes, and computing being done in absolutely every instant a processor is not occupied with communications, overall performance can be doubled or tripled on some applications, Lewis said. The problem is that programming the machine in this manner leaves it to the user to make sure that data are current and not changed while a message is being transmitted.

"Depending on how far you push it, the complexities go up," Lewis explained. In the Paragon, Intel has addressed the problem by installing  separate i860 microprocessor in each node to handle communications.

One of the most important performance metrics with MPPs is the ratio of the time the nodes spend processing data to the time they spend communicating with each other. If perhaps 10 or 20 percent of the time during  representative set of applications is spent communicating, the system is balanced. But much more than that indicates  dreaded communications bottleneck.

Cecil Leith,  computational scientist at Lawrence Livermore National Laboratory in Livermore, Calif., and a 40-year veteran of the computer industry, recently put a three-dimensional compressible hydrodynamics code (used in turbulence research) on a CM-5. Communications take up only about 15 percent of the time during execution, he



reported, but he frets that when the processors are outfitted with vector units and "get speeded up by a factor of 20, and communications by a factor of 2"—which is likely, in his view—"the program will become unbalanced again.

"The whole issue in parallel computing will be the relative speed of communications between the nodes to arithmetic on the nodes," Leith said.

Despite such concerns, he is an unabashed fan of the message-passing model. "All the synchronization is done implicitly, so all you have to know [when programming] is that you have the data you need before you start that phase of the calculation. It's similar to what many of us have been doing for decades" in moving data between mass storage devices and physical memory (a function handled by the operating system in virtual memory machines).

**BIG BANDWIDTH.** One important measure of the speed of internode communications is the so-called bisection bandwidth, the rate at which half the processors in the machine can send data to the other half. For most

competitive MPPs, this rate ranges from at least 100 megabytes per second for a machine with 20 or 30 nodes, up to about 12 gigabytes per second for a 1000-node one.

At least as important as bisection bandwidth is the latency—the time it takes for the processor in one node to prepare, or "set up," communications with another node, for the purpose of fetching data from the other node's memory, for example. In the CM-5, according to Leith, the latency is 350  $\mu$ s. After it gets started, message passing takes 0.5  $\mu$ s per word of data, but gathering the message in the source node and scattering it at the destination node can add another 3  $\mu$ s per word to the effective communications time.

"If you're going to send thousands of words, 350 microseconds doesn't matter much," Leith said. "But if you're sending 10 words, it matters a lot." A spokesman for Thinking Machines said that the 350- $\mu$ s figure was true of the company's preliminary system software for the CM-5, but has been improved in the current release to 110  $\mu$ s. Even so, most users expect latency to re-

main an issue with MPPs for the near future.

Because of latency, the more communications a program requires, the slower it will run on a message-passing architecture. Straightforward programs involving large vectors (arrays of numbers) will run well, but any application requiring frequent communications among nodes will bog down, according to David J. Kuck, director of the Center for Supercomputing Research and Development at the University of Illinois in Urbana.

At Boeing, for example, supercomputers—including MPPs—are used to simulate the electromagnetic signature and radar cross-section of military aircraft. Typically, the application requires solving an extremely large system of linear equations. This in turn requires generating a matrix of the coefficients of the linear system. This matrix may have a hundred thousand rows and columns.

In other types of applications involving large matrices, like computational fluid dynamics or structural engineering, the matrices are usually "sparse": almost all of the coefficients are zero. Not so in Boeing's

## Massively parallel computers at a glance

Company (date founded)	No. of machines sold to date	1991 sales, US \$ millions	Current model	Architecture	Node microprocessor	Peak execution rate per node, Mflops (double precision)	Comments
Intel Super- computer Sys- tems Division, Beaverton, Ore. (April 1984)	More than 325	90	Paragon XP/S	MIMD, 2-D mesh	Intel i860 XP (mesh)	50	Currently has one i860 XP per node for com- putations; next generation will have four i860s for computation
Kendall Square Research Inc., Waltham, Mass. (February 1986)	7	Not ap- plicable	KSR1	MIMD, hi- erarchy of rings	Custom super- scalar RISC chip	40	Uses hardware to present a shared-memory programming model, though memory is phys- ically distributed
MasPar Com- puter Corp., Sunnyvale, Calif. (March 1988)	More than 125	18.5	MP-1	SIMD, 2-D mesh, di- rect com- munication with 8 nearest neighbors	32 custom 4- bit processors per chip	1.2 (per 32 processors)	Has communication hardware integrated into processor chip (23 gigabytes per second for machine); through indirect addressing, processors have autonomy in memory ac- cesses
Meiko Scien- tific Corp., Waltham, Mass. (May 1985)	More than 425	25	Computing surface	MIMD, vari- able to- pology	Intel i860 or Sun Sparc	40 (with i860)	Links processing nodes with a Computing Sur- face Network, based on a custom-designed switching chip that can reconfigure the topol- ogy in response to software commands
nCube, Foster City, Calif. (1983)	More than 300	18	nCube 2S	MIMD, hypercube	64-bit custom scalar processor	2.4	Has integer and floating-point arithmetic, memory-control, and communications hardware integrated on processor chip
Parsytec GmbH, Aachen, Germany (October 1985)	Approx- imately 100	8	Parsytec GC	MIMD, 3-D mesh, vari- able to- pology	Transputer T9000	25 <sup>a</sup>	May be air- or water-cooled, depending on number of nodes
Thinking Machines Corp., Cam- bridge, Mass. (May 1983)	Approx- imately 90	85	CM-5	MIMD, fat tree	Sun Sparc	128 <sup>a</sup> (with max- imum of 4 vector units attached)	Has a time-sharing operating system; supports both MIMD and SIMD programming styles in hardware
Wavetracer Inc., Acton, Mass. (June 1988)	30	3	Data Trans- port Computer	SIMD, 3-D mesh	Custom- designed bit- serial processors	Not applicable	Works as an attached processor with Unix workstations or Apple personal computers

RISC = reduced-instruction-set computer MIMD/SIMD = multiple/single-instruction, multiple-data <sup>a</sup> Estimated performance; hardware not independently evaluated at press time

Source: Gary Smaby, Smaby Group Inc. (1991 sales); Jack Dongarra, Oak Ridge, Tenn. (peak execution rate per node); companies listed (all other information)



electromagnetics problem, where it is assumed that none are zero. This makes the task complex, because of the huge amounts of storage required. Yet these problems are tailor-made for MPPs, for several reasons.

The algorithms used to solve the equations have low communication-to-computation ratios, and the two tasks can be overlapped. Also, the task of generating the coefficient matrices is usually discounted, but in these applications on a vector supercomputer it is a significant bottleneck.

"People think of this as trivial," Lewis said. "But generating the system is something we cannot vectorize. Yet it is easily parallelized."

**SHARING MEMORY.** Because of the programming complexities of message-passing, several companies are attempting to develop MPP architectures that mimic the shared-memory model of mainframes and conventional (non-massively parallel) supercomputers. So far only one of these architectures—the KSR1 from Kendall Square Research—is on the market, but others from Cray Research, Convex Computer, Tera Computer, and other companies are under development.

The KSR1 can have up to 1088 nodes. Each has 32M bytes of memory, and although these memories are physically distributed, they are managed as if they were pieces of a single large virtual memory. The address of a certain data element, for example, is the same throughout the machine and regardless of which specific memory it happens to be in or which processor is trying to access it. When a processor needs the data at a certain address, the processor's own local memory is searched first. If the address is not there, a high-speed search engine finds the address and its data in another memory.

If the processor will later modify the data it needs, the address and data are transferred to the processor's local memory. But if the processor merely needs to read the data element, it is fetched and a copy remains in the memory that originally held it. Although the data element is moved about or duplicated, its virtual address remains the same, thanks to the search engine. "The view seen by a user, a compiler, or a third-party software writer is a sequentially consistent shared memory," said KSR's Burkhardt. All of the individual memories are managed like caches, on a first-in, first-out basis: a piece of data stays in a memory until newer data crowd it out.

The node processor is a 20-MHz super-scalar RISC chip, connected to coprocessors for communications with peripherals and for floating-point arithmetic. The floating-point chip has been rated at 15 megaflops in single-precision on the 100-by-100 Linpack, a common benchmark. There are no limits to the architecture's scalability, Burkhardt said, so, as with other MPPs, system performance is

constrained only by economics. "The only limit is volume or the amount of money you have, and you run out of money first," he noted.

Cray and Convex are also planning to adopt this type of memory strategy in their MPP architectures. Few details are available, however, since the machines are not expected to be introduced for at least a year.

Cray's MPP machine, code-named Tera 3D, will have memory that is "physically distributed and globally and logically shared," said Nelson, head of Cray's MPP project. The initial design will accommodate up to

## The next wave of massively parallel computers will give the illusion of shared memory

1024 nodes, each based on the Digital Equipment Corp.'s Alpha microprocessor, with a custom-designed memory controller in the form of a hardware shell around each Alpha to handle addressing (like the search engine in the KSR1). The design will make extensive use of Cray-developed technologies for the Y-MP and C90 lines, specifically for packaging, cooling, and input/output, Nelson said.

Cray's goal is a 150-gigaflops peak rate for a 1024-processor system, Nelson said. A second-generation MPP is to reach a teraflops (peak) in 1995, and a third is to sustain a teraflops rate two years later.

Convex Computer Corp., Richardson, Texas, is pursuing a similar approach, which it calls GSDVM, for globally shared, distributed virtual memory. But the company is not shooting for teraflops performance, according to staff mathematician Greg Astfalk. "Philosophically, our goal is to build a machine that is easy to use," Astfalk said. "We're not necessarily going after the highest level of performance. People-time [spent programming] is more expensive than machine time, so ease of use is important." Like Cray, Convex is pursuing a dual approach, developing in parallel both its next generation of vector supercomputers and its first MPP.

**DIVERSE AGENDA.** In recent years IBM has been pursuing an incredibly diverse agenda in supercomputing, including work on MPP architectures and systems at a laboratory in Kingston, N.Y. According to Irving Wladawsky-Berger, assistant general manager for supercomputing in IBM's Enterprise Systems business line, IBM is working on a distributed-memory, message-passing type of machine under a project code-named Vulcan. IBM expects to deliver variations of the

machine in the next year or two.

Computational nodes in the architecture are connected by a self-routing packet switch, according to Abraham Peled, head of IBM's power visualization product line, which is cooperating with Wladawsky-Berger's group in developing the machine. The scheme allows any node to be connected to any other via stages in the switch; the key to implementing the scheme was keeping the number of stages low so that the machine could be scaled up, Peled said. IBM engineers are striving to build an interconnection scheme with a latency, in internode communications, of less than 10  $\mu$ s.

Wladawsky-Berger believes the hardware and the software needed to effectively implement a shared memory model in an MPP with hundreds of nodes is still "a few years" away, but this is the direction he thinks the industry will take. In the machine of the future, in his view, "the memory will be physically distributed, but a combination of software and architecture—like in virtual memory—will make it possible for every node to address the whole memory. I think that's the direction everybody will take—that kind of hybrid."

One noteworthy company in this group is Tera Computer Co. of Seattle, Wash. Tera's machine will have up to 256 processors, each capable of processing up to 128 instruction streams simultaneously. The idea dates to the Heterogeneous Element Processor (HEP) computer, produced in the early 1980s by Denelcor, where Tera founder Burton Smith was vice president of research.

Tera engineer Robert Alverson likened the Tera machine's architecture to a dance hall, with "processors on one side, and memory on the other. But it's not a line dance: memory and processors are interspersed within a 3-D interconnection network. Each processor has memory near it, but a processor can also have its references spread out among all the memory in the system."

Each processor will have a peak rate of 400–650 megaflops, and sustain about 240 megaflops on matrix applications, Alverson said. There will be about 1G byte of memory per processor. Tera is shooting for an introduction in early 1994.

Boeing's Lewis is "intrigued" by the Tera machine because on certain benchmarks, he said, "I know I can't touch one of its processors with a 128-node" MPP of the current generation. "Of course, they can only provide results from a simulator at the moment," he added. And as with all previous multiple-instruction-stream machines—and all high-performance computers of any kind, for that matter—sophisticated compilers, which squeeze the most out of the machine while presenting a coherent programming model to the user, may be the single most critical factor in determining success or failure. ♦



# Software on the brink

*Operating systems and languages that will make supercomputers more widely usable are not far off*



First-time visitors to a foreign city that does not use their alphabet can feel totally disoriented, unless they spot those universal symbols for lodgings, restaurants, rest rooms, and telephones. Then hope springs anew that they will, after all, survive to explore this strange new place.

Similar signposts are being developed today for the world of supercomputing. The fundamental changes in architecture embodied in massively parallel supercomputing hardware at times confuse its developers, let alone users and programmers. Since the latter interact with a computer mainly through software—high-level languages and/or graphic user interfaces—software is key to making supercomputers less intimidating.

Software is certainly emphasized by supercomputer manufacturers. Kendall Square Research, Waltham, Mass., for one, has twice as many software as hardware engineers and Cray Research Inc., Eagan, Minn., now spends half of its R&D budget on software, or about 7 percent of its total 1991 revenues of US \$850 million.

Moreover, academic researchers have teamed up with manufacturers to adapt popular operating systems and conventional languages to supercomputing hardware. While parallel supercomputing has already spawned specialized environments and languages, for now most are adopting a "walk before you run" approach. By lending unconventional systems more of the look and feel of conventional ones, manufacturers hope to encourage scientists and engineers to apply supercomputers to solving extremely complex problems.

To do that, they are turning to Unix as a way of giving all kinds of supercomputers a common environment from which to run applications. The challenge is to agree upon a

Richard Comerford Senior Editor

parallel version of the popular operating system. For writing applications, academic and commercial technologists are rapidly hammering out a standard, parallel version of Fortran and trying to start down the same path for C. And work has also started on building tools for parallel programming that will work in any environment. The pace is frantic as workers try to make 1993 the year in which supercomputing, in the form of massively parallel processing (MPP), fulfills its promise.

**SOFTWARE PREMISES.** Supercomputing's basic goal is to attack complex problems that could never be solved in a reasonable amount of time before. From a hardware perspective, it is now much more practical to link 100 processors capable of 100 million floating-point operations per second than to build a 1-billion-floating-point-operation-per-second processor. So the current consensus is that parallel computing is the best route to top performance.

For software to make use of such an environment, it must be possible to divide a complex problem into many pieces, each of which can be worked on at the same time. While some problems must be solved one step at a time (what some refer to as a pathologically sequential data dependency), many industrial and scientific problems are not predominantly sequential. But the traditional von Neumann architecture dictated that programs be written that way.

Software that gives supercomputers the look and feel of conventional computers makes them easier for engineers to use

A simple parallelizable problem is a database search. Part of this search process is inherently sequential: call a record, search it for a bit pattern, note the occurrence of a match, note that the search of the record has been completed. But if the pattern being sought is given to multiple processors, they can search different records simultaneously, in parallel.

In scientific and engineering areas, problems commonly involve finding solutions to

sets of simultaneous differential equations using matrix multiplication. Since each member of one matrix is multiplied by one member of another, each multiplication can be handled by a different processor. With enough processors, many matrix multiplications can be performed at the same time.

**GRAND CHALLENGES.** Technologists have categorized various tasks they would like to take on by looking at the type of equations they involve and how much time it is practical to spend on them. (Even though the result may be accurate, a 72-hour-long program for predicting the weather 48 hours from when it starts running is not practical.)

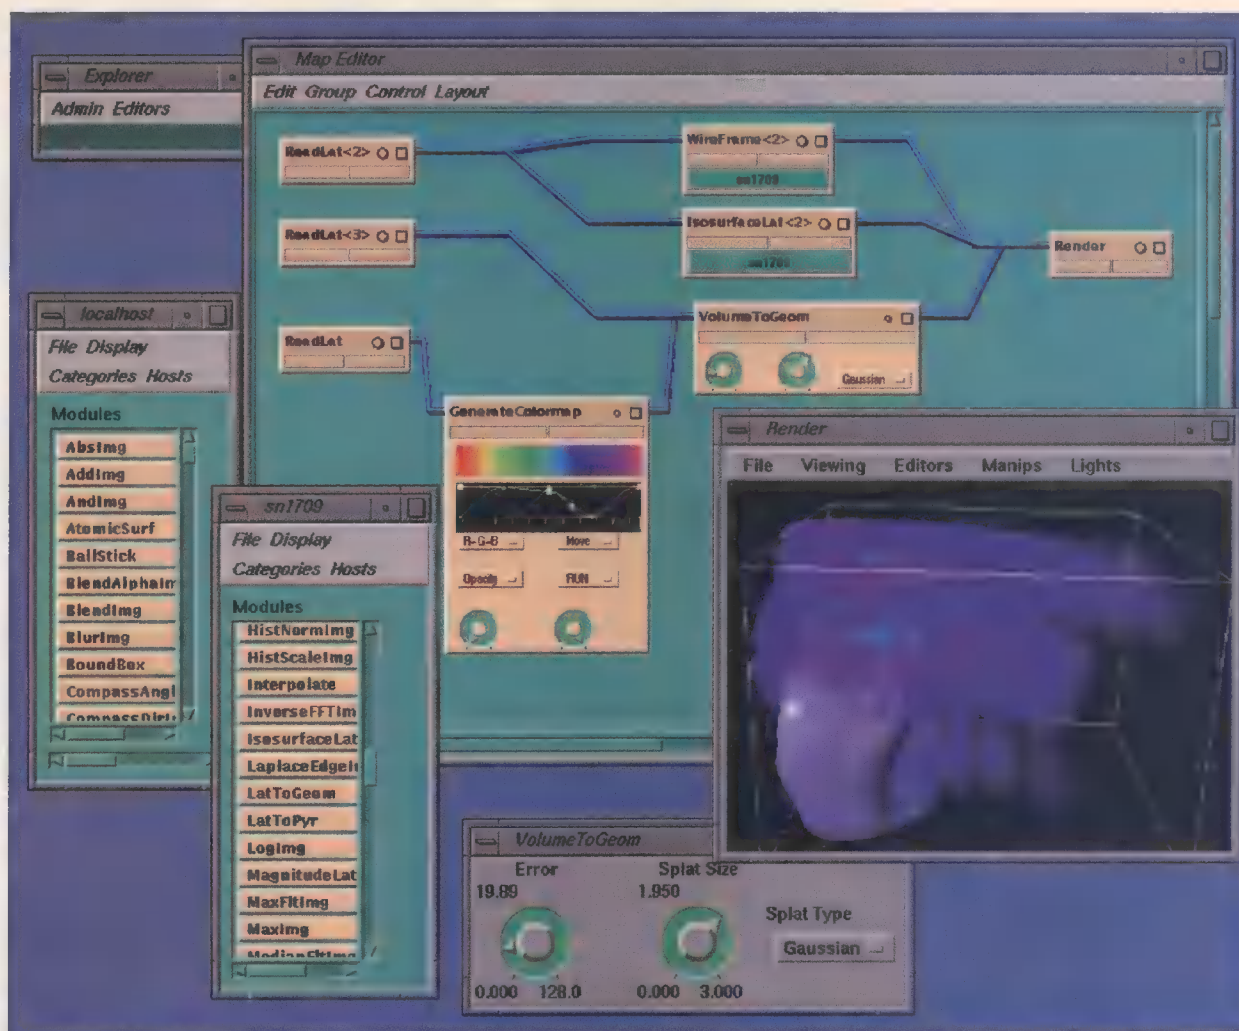
Combining such educated guesses with estimates of the rate at which the power of supercomputers is increasing, the U.S. Office of Science and Technology Policy's Committee on Physical, Mathematical, and Engineering Sciences correlates desired tasks and computing power in the report, *Grand Challenges: High Performance Computing and Communications*, issued in 1991 [Fig. 1].

What the report makes abundantly clear is that supercomputer hardware performance must soar by the end of the decade if the technical community is to succeed at their challenging tasks. So for work being done now to be useful in the future, software will have to be able to move easily up the gigaflops ladder.

As envisioned today, that will happen if operating systems and programming languages can be made to support multiple architectures. Much progress has been made recently, and some ongoing software work could bear fruit this year and early in 1993.

**UNIFIED OPERATIONS.** One after the other, supercomputer manufacturers are jumping on the Unix bandwagon, so that the operating system environment for supercomputing is quickly becoming more unified. At nCube, Foster City, Calif., a Unix look-alike is now running that Michael Meirer, chief executive officer of nCube, said comprises 85 percent of the standard Unix calls—built-in commands for standard functions. In its next release, nCube intends to support all Unix calls. "The industry is making tremendous progress toward defining a standard MPP model of Unix," he said. But Ted Tabloski, director of system software development at Thinking Machines Corp., Cambridge, Mass., feels that "Nobody [no standards organizations] at this point, with the exception





[1] Supercomputers promise to solve fundamental problems (Grand Challenges) involving large, complex arrays; what this requires in the way of computing capability has been mapped out [below right] by the Committee on Physical, Mathematical and Engineering Sciences of the U.S. Office of Science and Technology Policy.

For one of the problems—the conditions that result in a tornado-producing thunderstorm [above]—the simulation software is distributed between a Cray Y-MP and a Silicon Graphics Inc. IRIS workstation. The Y-MP creates the wireframe model, and the IRIS lets users interact with simulation variables.

## Defining terms

**Linear address space:** a conceptualization of memory system requiring that valid memory addresses be contiguous.

**Message passing:** any of a number of techniques by which processors and processes communicate status information in a parallel environment.

**Microcode:** the basic instruction set of a microprocessor implemented in its hardware.

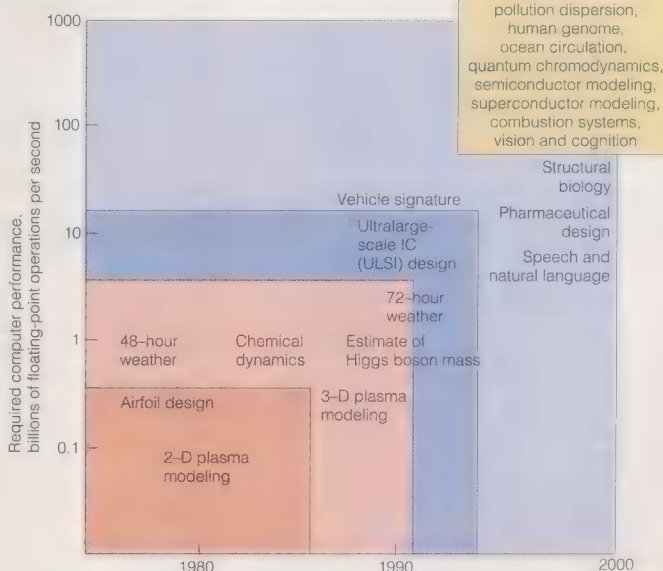
**Microkernel:** the core of an operating system containing elementary functions to which only higher levels of the operating system can gain access.

**Object:** a software construct consisting of a set of operations that share a state; the state consists of one or more values that can be viewed and modified only by those operations.

**Object orientation:** the property of a language that contains constructs by means of which programmers can easily create and manage objects.

**Partitioning:** the process of dividing software into segments that can be operated on independently and simultaneously.

## Grand Challenges



Source: U.S. Office of Science and Technology, 1991.



of OSF, is focusing on MPP operating system functions in their entirety."

In MPP as in workstations, there are two dominant Unix standards. One is OSF/1 from the Open Software Foundation Inc., Cambridge, Mass. OSF's research division worked with Intel Corp.'s Supercomputer Systems Division, Beaverton, Ore., to extend OSF/1 for Intel's Paragon massively parallel supercomputer. Kendall Square Research (KSR), the latest company to enter the massively parallel supercomputing arena, based its operating system, KSR OS, on OSF/1 with parallel extensions.

OSF/1's strongest competitor, Unix System V Release 4 (SVR4)—from AT&T subsidiary Unix Systems Laboratories in Summit, N.J.—is also hopping on the parallel-processing track. As part of the European Strategic Programme for Research and Development in Information Technology (Esprit), the ECU 14-million *Ouverture* program funded by the European Commis-

sion, Brussels, Belgium, aims to create a massively parallel version of Unix. A consortium that includes Alcatel-Alsthom, Chorus Systèmes, Olivetti, SGS-Thomson, Siemens-Nixdorf, Unix International, and Unix Systems Laboratories is weighing the technical aspects of integrating the Chorus/MiX microkernel architecture from Chorus Systèmes SA, Saint Quentin en Yvelines, France, into Unix Systems Labs' Unix SVR4.

Cray Research was the first supercomputing company to adopt Unix. In 1986 it released Unicos, which was based on Unix System V with Berkeley extensions. According to Bob Ewald, executive vice president of development, future versions of Unicos will let Cray unite massively parallel and vector systems of all sizes. The company's long-term thrust is to bridge the two hardware worlds with software. The company said it will have such a bridge available when it introduces its first MPP system in 1993.

When a bridge is built, applications will be

able to cross it to whichever type of machine best suits a given task. The modified Unicos will allow a Cray Y-MP C90, for example, to work "as one" with any of the firm's massively parallel offerings (the first of which will be based on the Alpha chip from Digital Equipment Corp., Maynard, Mass.). Compatibility with Posix and other standards will let applications move from single-chip Alpha workstations to Cray's massively parallel system.

That system is being designed so it can be coupled with a Y-MP; the operating system, user interface, and compiler will reside in a microkernel on the massively parallel system, which acts as an agent of Unicos. Which machine is performing the task will be transparent to the user, since the software parcels out the application to run at maximum efficiency.

**BIG AND LITTLE.** Thinking Machines also aspires to a world where the same programs can run on desktop machines and supercomputers. In its view, that benefit stems from having workstations and MPPs with extremely similar operating system software, as well as processors that share the same kind of architecture and microcode.

As Thinking Machine's director of software, Tabloski is especially proud of the firm's CM-5 operating system, which is based on Unix System V, too. A timesharing operating system, it was completed in 1989 and is still the only one, he said, for an MPP. (On the other hand, KSR supports the Tuxedo enterprise-wide transaction-processing system and Oracle 7 relational database management system, and others are reportedly working on delivering similar capabilities.)

The CM-5's control processors run a full version of Unix, with extensions for time sharing, networking, file systems, and parallel processing. The node processors run a Unix microkernel designed by Thinking Machines to support partitioning, message passing, and more so that the machine could switch contexts, protect user programs/privacy, and so on.

From a user's perspective, however, future versions of Unix-based operating systems may present not operating incompatibilities but simply choices based on performance and individual preferences. The IEEE's 1009, like Posix, aims to provide portability of applications to supercomputers running Unix-type software.

Further, Roel Pieper, president of Unix Systems Laboratories, and David Tory, president and chief executive officer of the Open Systems Foundation, announced on June 16 that the two groups would collaborate on bringing the two dominant forms of Unix closer. Said Tory, "It is important that we agree on standard specifications, and then implement products against those specifications." Unix Systems Labs' upcoming Unix SVR4.2 will support both the Open Look and Motif graphic user interfaces, and future releases will also support OSF's dis-

## Learning parallel programming

As future scientists and technologists, students must be taught how to program systems with parallel architectures. But when a supercomputer costs millions of dollars and is besieged by professors and graduate students doing important research projects, how does a university go about making it available to undergraduates? In most cases, the answer is they do not. Instead, institutions of higher learning are finding innovative ways of giving them hands-on experience.

Under a grant from the National Science Foundation, Washington, D.C., Chris Nevison, chairman of the computer science department at Colgate University, Hamilton, N.Y., has been conducting a workshop known as the undergraduate parallel computer (Uparc) group to establish hands-on parallel processing laboratories at numerous institutions. Uparc has devised a series of curriculum modules to instruct students in various aspects of parallel computing.

For instance, one module created by Daniel Hyde of the department of computer science at Bucknell University in Lewisburg, Pa., is being used by his students to set up a sort program with multiple communicating processes running on a single processor, and then to map that program to a multiple-processor environment. The module was tested by 15 students and, according to Hyde, it was very successful. The students "showed a genuine understanding of the [parallel processing] method," he reported.

Already, many laboratory course modules have been developed. The topics covered expose students to a wide range of problems that are often encountered in parallel systems. For instance, there are modules dealing with load balancing and scheduling, performance modeling, parallel searching, computing optimal binary search trees, converting sequential algorithms to parallel ones, and data flow.

During the summer, Nevison holds two-week courses for undergraduate faculty who can then bring the work to their students. The hardware used

in these laboratories is boards supplied by Computer Systems Architects (CSA) in Provo, Utah, which uses 32-bit SGS-Thomson/Inmos transputers. Popular in Europe, transputers are designed to work as message-passing processors in a parallel system. The message-passing system is based on the concept of communicating sequential processes developed by C.A.R. Hoare. With CSA's boards, the number of transputers may be expanded to as many as 1024; for his purposes, Nevison is using three stations, each having 17 transputers and a PC front end.

Occam and C are the programming languages for Nevison's courses. While learning a language is of value, it is more important in these courses to give the students an understanding of the concepts and principles underlying parallel computing. "Occam is very good for that," he stated. "A nice feature is its parallel construction; it is parallel from beginning to end." Occam lets students emulate parallel processing on a single transputer and then modify the existing code to run on a network of processors, meanwhile learning how messages can be routed efficiently.

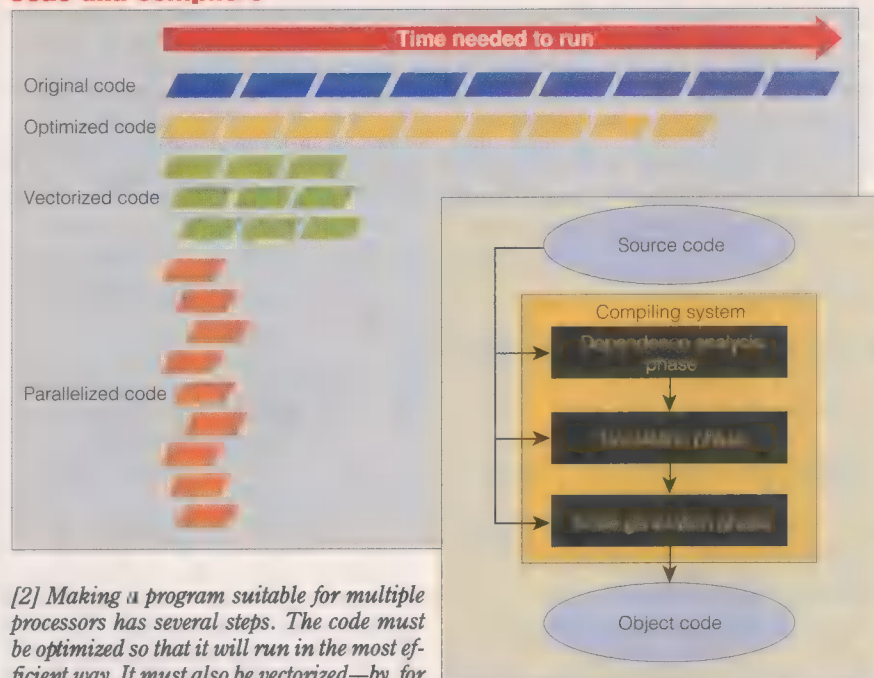
Occam was the first language designed to take advantage of the transputer's parallel-processing and message-passing capabilities. Today, CSA also supports Ada, C, C++, Fortran, Modula-2, Pascal, and Prolog parallel programming environments for the transputer; all use a concurrency library or language extensions to allow communication between multiple processors.

This month, CSA is sponsoring an intensive, two-day workshop to introduce teachers to the work already done in developing undergraduate parallel-processing laboratories using transputers. It is to be held Sept. 10-11 at the Alta, Utah, Lodge, and Nevison, along with Rod Tosten of Gettysburg College, Gettysburg, Pa., who has also been using transputers to teach parallel computing to undergraduates, will be among the instructors. More information about the workshop may be obtained by calling 801-374-2300.

—R.C.



## Code and compilers



[2] Making a program suitable for multiple processors has several steps. The code must be optimized so that it will run in the most efficient way. It must also be vectorized—by, for instance, changing a single scalar summation into multiple sums of its vector components. Finally, it must be parceled out so that it can run on multiple processors, simultaneously or with maximum overlap. Today's supercomputer compilers are capable of automating much of this process by analyzing code for independent elements that are parallelizable, translating it into parallel modules, and generating machine code. It is sometimes possible to make elements of such compilers reusable by building front ends to accept different languages and code generators to target different systems.

tributed computing environment (DCE), application environment specification (AES), and the OSF/1 functional specification.

**LANGUAGES AND DIALECTS.** With agreement on operating systems practically accomplished, "The single most important goal is to find a mechanism for machine-independent parallel programming," said Ken Kennedy, a professor at William Marshall Rice University's Center for Research on Parallel Computation in Houston, Texas. The programming model for most supercomputers relies on programmers supplying machine-dependent compiler directives to specify how data should be stored in memory, whether the data is private or shared, and where code is parallelized.

For performing scientific and engineering work on supercomputers, all manufacturers are sure that Fortran is essential. But at present, they are supporting several different versions of the language. For instance, Cray and nCube provide Fortran 77 compilers and are developing new ones for Fortran 90. Unlike Fortran 77, Fortran 90 is set up to efficiently handle large arrays of the sort routinely encountered in technical and scientific programming.

Thinking Machines, too, is a proponent of Fortran 90. The firm's marketing director, James Bailey, characterizes the '70s and '80s as "the decades of the 'DO loop,'" and the '90s as the decade of the array. In the future, as Bailey sees it, array math will go to

massively parallel machine while powerful workstations will run "dusty decks," the programs written for sequential processing (an allusion to the time when programs consisted of sets of 80-column Hollerith punch cards).

Kendall Square Research supports Fortran with Cray, IBM, or DEC extensions. But

The industry is making tremendous progress toward defining a standard MPP model of Unix

while Henry Burkhardt, its founder and chief executive officer, accepts that Fortran 90 will be an important dialect, he also thinks it does not suit a broad enough range of applications.

The general programming model of the massively parallel machine now under development at Cray Research has its roots in work done by Rice's Kennedy. Kennedy's group, working with Geoffrey Fox of the University of Syracuse, developed a parallel, or distributed, version of Fortran, called Fortran D, as part of research efforts at the university. Embodied in Fortran D is the no-

tion that decomposing a program's data into parallel streams implicitly decomposes the computations as well.

Some think that Fortran 90 is just the warm-up for High-Performance Fortran (HPF). This new version of the language is being worked on by approximately 35 entities—universities, national laboratories, and commercial companies such as Cray Research, Digital Equipment, IBM, Intel, and Thinking Machines. The group, known as the High-Performance Fortran Forum, had modest beginnings at a birds-of-a-feather meeting during the Supercomputing 1991 conference last November. With Kennedy in its chair, the forum held its first formal meeting at Rice University in January of this year; it drew 130 people.

**CLOSE AGREEMENT.** The group was in considerable harmony over what should be in the language. According to Dave Loveman, a member of the HPF forum and a senior consultant with DEC's Massively Parallel Systems Group, the forum's first goal is to "spell out what we agree on." He pointed out that HPF will include all of Fortran 90 and, following Kennedy's lead, will add compiler directives to help place and align data in memory so that parallel processors can divide work without conflicts. The forum plans to introduce HPF formally at the Supercomputing '92 conference at year-end. DEC and Intel have announced collaboration on an HPF compiler, which they ambitiously expect to have by the first quarter of 1993.

Ultimately, the forum's objective is to tailor Fortran that favors no particular parallel computer design and that can run efficiently on any system, regardless of its particular architecture (a feature that software designers refer to as being architecturally neutral or independent). The group means to offer the fruits of its work to all, but has no formal plans for standardization; that said, the co-chair of the ANSI Fortran 90 standards group, Maureen Hoffert of Hewlett-Packard Co., Fort Collins, Colo., is also a member of the HPF Forum.

Kennedy observed that the initial version of HPF will be just that: "It's not the whole solution but it's a good first step." He expected that, in a second round of activity, the language will be extended to handle sparse and irregular matrices. The compiler constructs needed to handle such matrices

are now understood in theory, he said, but "we haven't worked out all the details of implementation yet." He hopes that those details will be in place by 1994.

HPF should make development of applications for supercomputers attractive to a wide base of software vendors; the language will also be able to support applications for, say, networks of workstations. "You couldn't do it any other way," said Vincent Schuster, a member of the forum and president of the Portland Group in Wilsonville, Ore. (which has developed C and Fortran compilers for



the Intel i860). In the past, he noted, software developers had to redo their work every time they wanted to support a new system because "there was never a single programming model," a situation he believes the HPF Forum will rectify. For the Portland Group, HPF will permit it to also target processors like Sun's Sparc—"We won't be tied only to the 860," Schuster said.

**C, AND MORE.** While the new Fortran dialect may be the first widely supported parallel supercomputing language, other languages have their appeal. At IBM Israel Science and Technology in Haifa, Israel, researchers have developed compilers for a new, parallel version of C, called pC, that is an extension of ANSI C. ANSI not long ago accepted Thinking Machines' parallel version of C, called C\* (pronounced "see star"), as a base document for an industry-standard data-parallel C. Cray, Intel, Kendall Square Research, and others also support C.

All the same, parallel C trails Fortran 90, Thinking Machines' Bailey said, and other languages are still farther behind. "We think those [two] will be the languages that predominate," Bailey said. DEC's Loveman concurred that "C hasn't gone through the same standardization process that Fortran has." For instance, in going from the 77 to 90 standard version, Fortran acquired a consistent way of handling arrays; C has yet to move formally through such a procedure.

Today, the numerous versions of C being offered for parallel machines have all been customized by their vendors for those systems. Rice's Kennedy pointed out that several members of the HPF Forum want to start a standardizing effort for C also. He believes that the same techniques used for standardizing Fortran would be successful for C: "It's pretty natural to extend them to C," he told *IEEE Spectrum*. At Supercomputing '92, a birds-of-a-feather session being organized by Burt Holstead of DEC and Maya Gothale of the Supercomputing Research Center, Bowie, Md., will attempt to form a High-Performance C (HPC) group, which would hold meetings during 1993 at least.

DEC's Loveman noted that, when it comes to C, there are "lots of competing ideas," so there are some political issues in standardization. The technical challenge, he noted, is that an architectural view is heavily embedded in the C language. "C assumes a linear address space more than Fortran," he said, "It's more isolated in Fortran; in C, it's scattered throughout."

While this is no problem for machines such as Kendall Square's, whose Allcache architecture preserves the linear memory view for programmers, the goal of a parallel C would be to support all architectures, present and future. In the end, Loveman believes, a language must be stripped of any particular memory or architectural preconceptions, because they distract those using the language from finding solutions to their problems.

A more fundamental question that Loveman posed is: what do people actually want, C or C++? "C++ is a different animal," said Kennedy, observing of this object-oriented C-based language, that objects are a natural way to describe task or functional parallelism. An object orientation can simplify the definition, solution, and visualization of a problem, and so it is likely to close the distance quickly. Cray began delivering its compiler for C++ at the end of July and Kendall Square said its C++ compiler will be available in November. IBM Israel is working on C++ as is Tera Computer Co. in Seattle,

The goal is to  
create a Fortran that  
does not favor a particular  
computer design and can  
run on any system

Wash. Several university researchers, among them Manni Chandy at the California Institute of Technology in Pasadena and Dennis Gannon at Indiana University in Bloomington, are making it their business to create a standard version of parallel C++.

In view of the countless business applications that are inherently parallel—transaction processing, accounting, payroll, and database searching, to name a few—Cobol has its appeal. But it has so far seen "just a glint of interest, no one has really tackled it," claimed DEC's Loveman. Because of its memory architecture, Kendall Square is able to support Cobol, but a standard for parallel Cobol is nowhere in sight.

**NEW WAYS OF TALKING.** While many are attempting to make over traditional languages for nontraditional computer architectures, others believe that the new architectures would be better served by other kinds of languages. Take Sisal, a language created by researchers at Lawrence Livermore National Laboratory, Livermore, Calif. It is a functional language, one that expresses the nature of the tasks to be performed without regard for the underlying computer architecture. (Imperative languages, like Fortran and C, tell the system what to do in terms of its underlying subsystems.)

Until recently, one problem with functional languages has been building efficient compilers to turn the high-level descriptions into executable code. However, Convex Computer Corp., Richardson, Texas, working with Lawrence Livermore, demonstrated a Sisal compiler able to generate machine code about as fast as a Fortran compiler can; furthermore, their resultant code was even comparable in performance and memory requirements. As massive parallelism becomes a more established architecture, functional

languages will catch on; for now, they are likely to remain a research activity.

**TOOLS.** Standard operating systems and languages/compilers are two of the three legs needed to support supercomputing; the third is the tools for application development, which are just now becoming a focus—particularly those that allow programmers to visualize performance.

Having created the first supercomputer, Cray Research has been working on the problems of supercomputer programming longer than any other firm. It has also spent more time developing and refining software support tools than any other vendor, according to Mark Furtney, a group leader at Cray Research responsible for tools.

One of Cray's tools is atexpert, an expert system for what the company calls autotasking—the automatic division of a program into several parallel processes. Realizing that time on supercomputers is often at a premium, programmers resort to atexpert, which uses measurements taken on a production, rather than dedicated, system, to see how effectively their programs run in parallel. Using an X-Window system

interface to graphically display data, atexpert shows where a program is spending most of its time and whether those areas are being executed serially or in parallel. The software analysis system also suggests ways in which performance could be improved.

A particularly useful forthcoming tool is a program browser called xbrowse, for use in viewing and editing Fortran codes. Depending upon what code is being reviewed, the language-sensitive browser provides information on the routine, file, or the whole program level; even if a program is made up of several files, the browser treats it as a single unit. The browser can also help the user insert directives in the code if they are needed to help the compiler distribute work. For example, the browser will help transform loops into parallel operations.


Cray plans to make similar tools available for its MPP supercomputer and is looking at the possibility of supporting other manufacturers' systems as well. In the past, should a programmer need to switch to a different system, he or she would have to become familiar with another set of tools and tool interfaces.

Currently, Dan Reed and his students at the University of Illinois at Urbana are developing machine-independent performance-visualization tools. The group has already developed a portable visualization system that can be used with any language (provided the compiler can insert appropriate event logging) and is collaborating with researchers at Rice University to find ways to use this system with Fortran D and HPF. Just as standards are dawning for parallel operating systems and programming languages, so will standard computer-aided software engineering tools for supercomputing soon see the light of day.



# Revolution or evolution?

*The best approach to massive parallelism is a holistic one that fully incorporates software requirements*



The key question about parallel supercomputing is not whether it will happen, but how it will evolve. The rise in the fastest machine clock speeds over the past few decades finally seems to be leveling off as designers begin feeling the constraints of fundamental physical limitations imposed by the speed of light. To avoid this stalemate, many are looking to parallel machines, in which many processors execute independent computational operations simultaneously.

The concept of parallel computing is not new by any means: such pioneers as Charles Babbage, Lewis Fry Richardson, and John von Neumann all had ideas on the subject. Indeed, the current revolution in parallelism arises not from any conceptual breakthrough but from our new-found abilities to realize parallel systems in hardware. Thus the ideas are old, the technology is new, and the software is relatively unexplored. The greatest Grand Challenge is to explore these implementations and understand how they fit users' needs and nature's constraints.

With this in mind, consider two approaches to achieving practical thousandfold parallelism by the year 2000. One approach, which we label *revolutionary*, calls for immediately building machines with thousands of processors, using whatever architectural ideas are currently deemed feasible from the system packaging viewpoint. Eventually, it is hoped, the systems software and application program base will catch up.

The other, *evolutionary*, approach is to methodically double the number of processors each year, while nurturing and enlarging the software base correspondingly. A starting point of, for example, 16 to 32 processors this year promises to become well over a thousand processors in 2000.

In both cases, the end goal is to have general-purpose large-scale parallel supercomputers by the year 2000. Which approach is better? Without a doubt, the evolutionary one.

Nonetheless, a revolution is being fomented by almost all the newer players in commercial supercomputing and most of the U.S. government agencies participating in the Federal High Performance Computing and Communications Initiative. The justification for this preference hinges largely upon peak theoretical speeds. Massively parallel machines can be built with impressive theoretical performance, and it is believed that committed users will eventually learn how to harness much of that theoretical speed.

Later on, it is said, supporting software systems will catch up, allowing more users—including less experienced ones—to gain access to those machines' power. It is argued that vector supercomputing came about that way, though we would counter that parallel computing is significantly more complex. This point of view also neglects the critical fact that the Cray-1's scalar performance was the world's fastest when the machine was introduced in 1976, whereas a modern massively parallel machine's single-processor speed is typically equivalent to that of a decent workstation.

The arguments for an evolution are that

Ideas on massive parallelism are old, the technology is new, and the software is unexplored

parallel architecture and software are very complex and poorly understood, so they must evolve together. Peak theoretical speed matters much less than delivered performance on real application programs. Users are more likely to embrace a gradually evolving software base than to radically reimplement or rewrite all the code they have. Furthermore, the results of the process may not be a single system: evolution is naturally suited to branching toward mul-

multiple systems for multiple types of computation, as has been true in sequential computing.

The revolutionary approach presents both users and systems software designers with a moving target. For example, the radical architectural changes between the Connection Machine 2 (CM-2) and CM-5, both from Thinking Machines Corp., Cambridge, Mass., strongly suggest that massively parallel computer designers are still experimenting with their architectures. Such discontinuities in system software and architecture tend to discredit the vendors and discourage the consumers.

These issues are far from academic. National industrial competitiveness depends on a strong, high-performance computing base that can support technological and industrial design processes. Those processes will need sustained, improved support in the coming years, not vague promises of future miracles.

**BIRTH OF AN ARCHITECTURE.** Manufacturers began trying to deliver on the promise of massive parallelism in the mid-1980s, when the first commercial machines of this type hit the market. This development had two causes. One was the increasing difficulty of deriving faster supercomputers from faster circuits alone. The other was the wide availability of rather fast and rather inexpensive microprocessors.

The 1980s was a decade of stupendous ferment in computer systems: almost every architecture then imaginable was built or attempted, and incredible performance claims were routinely made and accepted. Thus the expectations of those who use powerful computers are very high. By now, the practicality of massively parallel systems with teraflops performance—a trillion floating-point operations per second—is an article of faith with many people. But there is regrettably little discussion of what is actually meant by teraflops performance, and little evidence that ordinary users will achieve it on real codes in the foreseeable future.

In short, parallel machines have always demonstrated peak speeds that are very high for their era, and their software has uniformly been of low quality [Table 1]. What has changed is that the hardware technology has caught up with the field, making it easy for companies to produce such machines now.

Consequently, the 1990s appear to present a parallel processing imperative: push

George Cybenko Dartmouth College  
David J. Kuck University of Illinois



ahead with the revolutionary approach to parallel computing, or risk falling behind in those areas of research and development that increasingly depend on high-performance computing. But the validity at this time of the imperative, in free-market economic terms, is debatable. For those contemplating buying a massively parallel computer, the key questions are: when should I make a serious investment in parallel computing? And when will parallel software catch up?

Often, managers find it extremely difficult to evaluate the arguments of engineers and scientists for purchasing a parallel system. Technical people want their systems to stay competitive in performance.

The best technical people have always risen to extreme challenges and will work hard to impress management as well as their peers. This effort can lead to great economic success for system vendors but to various "false successes" for the purchaser in an economic or technical sense. Among the common scenarios:

- The programs developed lack the functionality or utility of those they are intended to replace, and so are generally regarded as useless.
- The programs developed fail to run much faster than they did on previous equipment; sometimes the new, parallel programs will

run just as quickly on the multiple processors of a conventional, vector supercomputer as on the massively parallel machine.

- The programs developed are useful but required a great deal of time. Furthermore, little was learned in the process that can be reused on the next problem.

It seems inevitable that the 1990s must be a decade of convergence in parallel processing. The market cannot tolerate large discontinuities in performance, architecture, and programming models. Scientists and engineers now expect teraflops machines by 1996—and in a climate of heightened expectations, four years is a long time.

Price estimates of these first teraflops machines, based on hardware and a modest profit margin, exceed US \$100 million. If experiences with such machines are like those outlined above, the field of massively parallel computing risks taking a spot on the technical community's back burner, much as artificial intelligence did after overextending itself in the 1980s.

So the collective goal of the community should be convergence on usable high-performance parallel systems, and "usable" implies easy programmability, program portability between systems, and stability of performance from one application to the next, from one machine to the next. In the long run, parallel processing will surely prevail,

but crucial architectural issues remain to be settled, and many software issues remain to be resolved. It is only a matter of time before parallel computing offers high performance in a consistent, dependable, and cost-effective way.

Two important technical points are implicit in the preceding discussion. First, we need cost-effective programming techniques and tools, such as good compilers for established languages. Second, performance must be delivered to a wide range of users—an obvious need is to collect and disseminate performance information and to improve systems on the basis of this information.

**MEASURING PERFORMANCE.** During the 1980s, steady progress was made in evaluating the performance of conventional (vector) supercomputers. We started with peak speeds, went through Livermore Fortran Loops and Linpack benchmarks during the 1980s, and began the 1990s with the Perfect and Numerical Aerodynamic Simulation (NAS) Parallel Benchmarks. The progression is clear. What counts is performance delivered on user codes and not idealized computation rates based on the capabilities of exotic semiconductors. Unfortunately, there is now a risk of lapsing back into that earlier era, because measuring progress by peak theoretical speeds is both easier to accomplish architecturally and more impressive to the onlooker. (Peak speeds have been cynically defined as the "speed of a computer when not running any software.")

What then is the reality of current supercomputer performance? A study at one of the National Science Foundation supercomputer centers has found that the average delivered performance on the center's Cray Y-MP has been less than 25 percent of its peak theoretical single-processor performance (75 megaflops out of a possible 333). That figure is quite impressive for a single-processor supercomputer, given the wide array of users and applications [see also p. 60]. It is probably slightly lower than would be found at national laboratories but above that at many industrial sites, where supercomputers tend to run software written by third parties.

The Perfect Benchmarks were developed by a consortium of supercomputer users and vendors and coordinated by the University of Illinois, Urbana-Champaign. The benchmarks were designed to measure the performance of the overall system, including compilers, on small problems representative of actual supercomputer applications. Using compiler optimizations only, that is, without human intervention, the benchmarks execute at about 1 percent of the theoretical parallel peak speed on all current conventional supercomputers.

Taken together, these results suggest that the typical user of a conventional supercomputer should expect somewhere between 1 percent and 25 percent of the peak speeds. Higher rates can be achieved in situations with tuned programs running on very large

## 1. A brief history of massively parallel computers

Decade	Machines	Relative to decade:		
		Peak speed	Hardware technology level	Software quality
1960s	Illiac IV (University of Illinois) PEPE (Bell Laboratories)	High	High	Low
1970s	Staran (Goodyear) DAP (ICL)	High	Medium	Low
1980s	Hypercubes SIMDs	High	Low	Low

SIMD = single-instruction, multiple-data stream computer

## 2. Vector and massively parallel supercomputers compared

No. of processors	System	Peak speed, gigaflops	Benchmark code speed, gigaflops	Efficiency (code/peak)
NAS Parallel Benchmark EP (size: 2 <sup>28</sup> )				
16	Cray C90	16	8.3	50%
128	Intel iPSC/860	2.6	0.4	15%
65536	Thinking Machines CM-200	20	2.4	12%
NAS Parallel Benchmark FFT (size: 256 <sup>2</sup> × 128)				
16	Cray C90	16	4.5	30%
128	Intel iPSC/860	2.6	0.5	20%
65536	Thinking Machines CM-2	14	0.5	4%
NAS Parallel Benchmark CG (size: 2 × 10 <sup>6</sup> )				
16	Cray C90	16	2.6	16%
128	Intel iPSC/860	2.6	0.07	3%
16384	Thinking Machines CM-2	35	0.1	3%

Source: NASA Ames Research Center, Technical Report RNR-92-002

NAS = Numerical Aerodynamic Simulation

EP, FFT, CG: text



homogeneous sets of data. It can be argued that such programs are rare in industrial settings but do occur in such areas as quantum chromodynamics (the study of the interactions among quarks and gluons in subatomic particles).

How do massively parallel supercomputers compare with traditional supercomputers on codes tuned by experts? Some government labs have tried to answer this question in order to assess the potential of massively parallel machines for their needs.

At the National Aeronautics and Space Administration's Ames Research Center, Moffett Field, Calif., for example, some pseudo-applications have been designed to be representative of the type of fluid dynamics and aerodynamics simulations conducted there [Table 2]. These pseudo-applications are known as the NAS Parallel Benchmarks. Three of the codes used were EP, FFT, and CG, which are, respectively, a highly parallel Monte Carlo simulation, a three-dimensional Poisson partial differential equation solver using fast Fourier transforms, and a conjugate gradient linear equation solver for a banded system of equations. The codes were all optimized manually and individually for each machine tested.

The results show that on these codes, only the Cray currently demonstrates gigaflops performance. Thus, in relation to complete application programs running with problems of various sizes, massively parallel computers have efficiencies well below 1 percent.

Performance information about complete applications is harder to obtain because porting large complex programs to massively parallel machines has proved rough going. Often, after intensively rewriting an application code for a massively parallel machine, researchers have found the parallel performance to outdo the performance of the original traditional supercomputer code.

A group of investigators at Los Alamos National Laboratory in New Mexico undertook the explicit parallelization of a few key applications commonly used at that facility. Again, the performance of the massively parallel codes when run on a parallel machine overshadowed the performance of the original vector codes. However, the Los Alamos group then took the rewritten code and reran it on the Cray, a Y-MP/8. They found that the tuning and optimizations that went into parallelizing the code for a data parallel machine improved performance on the eight-processor Cray as well. For all three codes tested this way, the Cray far outdid the massively parallel machine.

As these examples show, the true performance story behind massive parallelism is hard to come by and not entirely flattering to the architectures. Realizing good performance on real application programs is a hard problem—one that has yet to be solved by designers of massively parallel machines. The challenge before those designers is to

evolve their machines toward architectures that support high sustained computational rates on real programs. This advance must lead to a convergence of designs in the 1990s.

**GOING INTERNATIONAL.** Mathematical equations govern many components and aspects of computers, but for overall systems we lack even generally accepted performance units, apart from elapsed time. In a sense, we are like primitive people who could count days by moving stones among holes, and could deal with earth, fire, water, and wind independently, but had no grasp how the

'The true performance story behind massive parallelism is hard to come by and not entirely flattering'

world worked, much less how to control or improve it.

Take airplanes. They are complex systems that can be classified by cruising speed and range, passenger and freight capacity, fuel consumption and waste emission per mile, runway length requirements, and so on. These figures allow airlines to base purchasing decisions on fairly clear parameters.

With supercomputers, on the other hand, we can only discuss peak speed and how long they take to do specific computations. Using the airplane analogy, it is as if we needed to quote the plane's characteristics for each flight plan. So we may be forewarned that a certain plane is great between Chicago and Los Angeles, will not be competitive between Chicago and Boston, and concerning other pairs of cities the buyers should check the details.

At this point, though, we are concerned less with consumers of parallel supercomputers than with their designers. How can we solve the greatest Grand Challenge and produce a breakthrough parallel system that performs well on all scientific and engineering computations? Or, alternatively, how can we determine that this is not practical today (or possibly ever) and then take the appropriate steps?

Many benefits would flow from measuring, reporting, studying, and steadily improving system performance on real applications. The undertaking could be based on a comprehensive performance database that would be available on a national network: a national performance center. Everyone supports this concept in the abstract, but no action has been taken. It is a huge undertaking that requires the support of Government, companies, and users alike.

Most academics see it as too much work

for too little glory and seem to prefer to start some other effort that they can claim as their own. Government agencies, too, seem to react negatively if the numbers obtained fail to support their individual agendas. Industry has supported the Perfect benchmarks for high-performance computing, and vendors formed the Simulation of Propulsion Engine Cycle (SPEC) benchmarking consortium to analyze and compare workstation performance. Now, a project on a truly grand scale is required, and there is no agreement on how to do it or who will pay for it.

The fact that industry has broadly cooperated in Perfect and SPEC indicates that it is really concerned and understands the problem much better than those in Government or academia. However, the cost must in part be borne by Government, and the work must be in part be done by academia if the project is to have the breadth essential to success. In fact, the performance meta-center would integrate the otherwise disparate computing and communication aspects of the U.S. government's High Performance Computing and Communications Initiative (HPCCI).

The crucial importance of a gigabit-per-second network would be clearly demonstrated as distributed performance data, benchmark, and Grand Challenge codes from across the country were made available to the community.

Regardless of whether the best path to practical parallel processing is by evolution or revolution, it is essential to have the quantitative basis for decision-making that would be provided by a national performance metacenter. Our HPCCI goals would be clarified by defining the database contents, and progress toward these goals would be immediately spread throughout the HPCCI community. The approach would also crystallize agendas for future action in the research and engineering communities: compilers and numerical libraries would be improved, new architectures would have to deliver performance—not just show peak-performance improvements. Users would have broad-based comparable information on which to base their choice of computing engine. Finally, administrators would have rational grounds (which do not exist today) for observing and measuring progress.

This database would give rise to a feedback mechanism that could help to correct most of the problems apparent in the HPCCI today. This is an urgent need and an absolutely necessary condition if we are to achieve near-term success in this program.


**ABOUT THE AUTHORS.** George Cybenko [SM] is the Dorothy and Walter Gramm Professor of Engineering at the Thayer School of Engineering, Dartmouth College, Hanover, N.H.

David J. Kuck [F] is professor of computer science and electrical engineering at the University of Illinois at Urbana-Champaign and director of the university's Center for Supercomputing Research and Development. ♦



# Japan: a competitive assessment

*The Japanese supercomputer industry is a formidable challenger—lagging the United States in some areas, leading in others*



Home to a quarter of the world's 500 or so vector supercomputers, Japan gets high marks for both effort and achievement in all aspects of high-performance computing. The country funds research in information technology more heavily than all other fields except the life sciences and the environment, and has the only *bona fide* supercomputer industry outside the United States.

Comparisons of the two countries' accomplishments and approaches to supercomputing are inevitable, ■ are the periodic claims, by manufacturers of either country, of having the world's most powerful machine. Japanese supercomputers are built by the same vertically integrated giant companies that produce all kinds of technological products, from cellular telephones to semiconductors. So their approach is hardly surprising: extremely fast individual supercomputer processing units built around very high-speed emitter-coupled-logic (ECL) semiconductors of their own fabrication, in which they are second to none.

As in the United States, the information technology research establishment in Japan has industrial, governmental, private, and academic components. In 1989 (the most recent year for which figures are readily available), the total R&D budget was 1012 billion yen (about US \$7.5 billion), according to *Indicators of Science and Technology*, published last year in Tokyo by Okurasho. Of this total, 958 billion were spent by industry, 24 billion by private research institutes, 23 billion by universities, and 5 billion by governmental research institutes. (There were about 133.5 yen to the U.S. dollar in 1989.)

**DEEDS AND DESIGNS.** Japan's big three in supercomputers are NEC Corp., Fujitsu Ltd., and Hitachi Ltd., all headquartered in Tokyo.

David K. Kahaner U.S. Office of Naval Research Asia  
Ulrich Wattenberg  
German National Center for Computer Science

About half of the 125 vector supercomputers in the country are Fujitsu models; Hitachi, NEC, and Eagan, Minn.-based Cray Research Inc. share the balance.

Japanese manufacturers always derive peak performance figures from the hardware specifications. It varies from 5 billion floating-point operations per second (gigaflops) for the Fujitsu VP2600 to 32 gigaflops for Hitachi's recently introduced S-3800; NEC's SX-3 has a peak of 26 gigaflops. For comparison, Cray's Y-MP C90 has a peak of about 16 gigaflops.

These figures, however, are derived merely by multiplying the number of processors in the machine by their peak rate. Real applications run far more slowly. Informally, many scientists assume that usable speed is one order of magnitude less than claimed peak; this sustained rate is heavily influenced by how rapidly and in what quantity data can be moved around, how much of the calculation is array-based, how often arithmetic processing is interrupted by comparative operations, and other factors.

Clearly, Japanese supercomputers have very fast central processing units (CPUs). But at the time of writing, only one company—NEC—has ■ multiprocessor system installed outside its own facilities. NEC's single processors have performed outstandingly, but not the multiple-CPU system (which has fallen very short of its potential peak, even discounting an order of magnitude). At the moment NEC and the other Japanese vendors are several years behind Cray, although as they gain experience they are bound to narrow the difference.

Various supercomputer benchmarks, which were intended to be better indicators of performance than the theoretical numbers, have been bitterly criticized as oversimplified and incapable of reflecting realistic complexity. They do indisputably show, nonetheless, that the Japanese single-CPU systems have vastly improved, to the point where they are internationally competitive.

**FASTER COMPONENTS.** Of course, the trend today is toward parallelism, and a distinction must be drawn between single- and multiple-CPU performance. In practice, high performance is much harder to achieve on multiple processors, since it demands more balance in both hardware and system software and autotasking capabilities. In general, supercomputer designers try for an equilibrium among memory speed and capacity, arithmetic processor performance, data movement capability, and so on. This has

been one of Cray's particular strengths.

Japanese manufacturers, on the other hand, have generally taken another path toward high performance. Their approach was perhaps best summed up by NEC in its 1990 annual report: "NEC's approach to supercomputer architecture is clear. Our first priority is to provide high-speed single processor systems which have vector processing functions and are driven by the fastest technologies, while giving due consideration to ease of programming and ease of use." All within the limits of economic feasibility, of course.

From this it follows that performance, for the Japanese machines, depends on four hardware goals: faster chips, smaller size, heat reduction, and elimination of logic bugs. Supercomputers from NEC, Fujitsu, and Hitachi use tried-and-true ECL semiconductors for basic processor chips, but have pushed their capabilities. For example, clock cycle times are 3.2 ns (Fujitsu), 2.5 ns

Hitachi Ltd.





(NEC), and 2.0  $\mu$ s (Hitachi). These figures are better than the ECL in the Cray Y-MP C90, with its 4.0-ns cycle time. Although clock rate is only one factor in overall system performance, it is an important one.

Behind these impressive clock rates are equally impressive feature densities on integrated circuits. ECL gate densities are on the rise; the supercomputer announced by Hitachi earlier this year uses 25 000-gate arrays, for example. Since 1989, NEC has used 20 000-gate arrays, and Fujitsu, 15 000-gate arrays. At the high end, all these machines have water-cooled CPUs, but slower, air-cooled versions are also available. In addition, air cooling is used in peripheral devices.

So far, gallium arsenide and Josephson junctions have not been employed for CPU chips in any commercial Japanese machine. Fujitsu uses GaAs chips in some of its peripherals so these can be effectively cooled by air. But by and large, the Japanese have held back from commercializing exotic IC technologies, the prevailing opinion being that more performance can still be squeezed out of silicon. But the Japanese do see GaAs as slowly replacing silicon ECL, and have plenty of research projects going on in preparation for this transition.

Fujitsu, for example, has developed a very large-scale integrated (VLSI) Josephson-junction chip, and plans to use it in its next-

generation supercomputers, probably out in the mid-1990s (development of a high-end supercomputer takes three to five years). Similarly, NEC has developed GaAs logic and memory devices and a GaAs multichip package for supercomputers. Last, but far from least, NEC recently demonstrated a 64-bit GaAs processor, 17 mm on a side and air-cooled, with a peak performance of all of 0.2 gigaflops.


The Hitachi S-3800 has the shortest clock cycle of any commercial computer in the world; it is two billionths of a second

In terms of CPU architectures, one of the principal technical differences between U.S. and Japanese supercomputer CPUs is the number of pipelines, a standard feature of high-performance processors. In a pipeline, several instructions are in various stages of execution at the same time: an assembly-line approach to executing instructions. Only part of the processing is done at each stage, and the passing of an instruction through all

the pipeline's stages completes its execution. (CPU pipes have only one instruction path and must all carry out the same calculation at the same stage, whereas different instructions can be executing on the independent CPUs.)

U.S. machines have more CPUs, each of which has a small number of pipes. Japanese machines have fewer CPUs but each has more pipes—as many as 16. This distinction is mostly due to greater U.S. experience in building multi-CPU machines, and is slowly changing as the Japanese add more CPUs to their systems.

**ACADEMIC ACCESS.** Most university scientists in Japan can get supercomputer time, but rarely on top-end machines, which are mainly to be found at industrial laboratories or in the prestigious national universities. Even so, this is an improvement on the state of affairs two or three years ago, though it is still below what is available to U.S. academics. There is nothing comparable to the four supercomputer centers in the United States partially sponsored by the National Science Foundation, even though supercomputer centers exist at leading Japanese universities and at several government and some corporate laboratories. The Recruit Group's Institute for Supercomputing Research and the Institute of Computational Fluid Dynamics are good examples of the last named.



Introduced last April, Hitachi's S-3800 supercomputer has up to four processors, each of which is said to peak at 8 billion floating-point operations per second (gigaflops). The 32-gigaflops aggregate peak is the highest claimed for any vector processor.



## How Japan's 'Big Three' stack up

Company	Current supercomputer family	Claimed peak execution rate per CPU; maximum number of CPUs	Central processing unit (CPU) cycle time, nanoseconds	Operating systems	Comments/special features
Hitachi Ltd.	HITAC S-3800	8 gigaflops; 4 CPUs	2	Unix, VOS 3/AS	Scientific Animation Graphics allows direct access to high-speed storage and can send out ordinary (NTSC) or high-definition TV signals at 30 frames per second
NEC Corp.	SX-3R	6.4 gigaflops; 4 CPUs	2.5	Unix	Has been aggressively marketing SX-3 line in the United States; a few have been installed outside Japan and performance measurements are starting to appear
Fujitsu Ltd.	VP2000	5 gigaflops; 2 CPUs	3.2	Unix, MSP-EX	Each pipeline in the CPU can perform up to 16 floating-point operations per clock cycle; a new model, the VPX, has essentially the same performance but runs only Unix; has stepped up marketing efforts in the United States, but is aiming at private industry sales only

About 40 supercomputers are installed at various Japanese universities, but at least a third are older machines or others with very modest performance. Still others have non-standard operating systems, few standard application software products, and inadequate networking. These last might still be useful, of course, for training and non-demanding applications.

Supercomputer networking in Japan has improved of late, but there are more high-performance networks in the United States than in Japan. U.S. network interconnectivity is also much better, although several more or less independent Japanese networks are supported by different government ministries.

Counterparts to the very high-performance networking projects in progress or planned in the United States have not yet jelled in Japan. However, the country has excellent—and in some cases, unique—technology, including a large infrastructure in integrated-services digital networks (ISDNs), and any obstacles to networking there seem to be more social, organizational, or cultural than technological.

As for academic supercomputer networks, they are less ubiquitous than in the United States. The prestigious Japanese universities have excellent services, but many lesser universities have none at all. In general, research in supercomputing lags behind that in the West, except for application developers working on commercial software packages. One or two supercomputer conferences are held each year, and at their small technical programs, fewer than one-third as many papers are presented as at a good U.S. conference.

**UNIX RULES.** Software for supercomputers includes compilers, libraries, operating systems, support for networking, and software tools. All three Japanese supercomputers now are available with a version of the Unix operating system, which will help the migration of applications onto Japanese systems. The industry is only now coming to grips with the need to assess software costs, and moving to Unix is seen as one way to reduce

costs for end-user and vendor alike.

This switch away from the use of proprietary operating systems has occurred only in the past two or three years. For Hitachi it is only just now occurring, and the company has not totally embraced Unix; its new supercomputer is available in a Unix version, and also with the company's own IBM-like operating system, for compatibility with older Hitachi systems. Fujitsu likewise supports both Unix and its own operating system.

In the past, applications developed in the West have been installed very slowly, which impeded sales of Japanese supercomputers both inside and outside their native land. Using Unix will ameliorate this situation, to the extent that it guarantees greater software portability and shorter development time. Efficient program execution, of course, is another matter entirely. As yet no shortcut to maximum performance seems to exist besides incorporating knowledge of the hardware into the algorithms and software.

In the early days, Japan's development of supercomputer software was limited to

The Japanese produce first-rate applications for supercomputers, but in software in general are behind the United States

producing Japanese language interfaces for Western software products. This activity is still important. Recently, for example, NEC interfaced the latest version of Nastran, the heavily used engineering analysis system, to its machines. The company's promotional literature lists about one hundred software products (many from the West), and Fujitsu and Hitachi are engaged in similar projects.

But by now, first-rate packages designed

and implemented in Japan are appearing. Good examples include Hitachi's Deqsol (differential equation solver), for the solution of the partial differential equations arising in engineering simulation; Alpha-flow from the Fuji Research Institute, for solution of fluid dynamics problems; Fujitsu's Fortran/KR for allowing object-oriented programming from within a Fortran environment; and Amoss (for ab-initio molecular orbital system) from NEC for molecular orbital calculations.

There is also a trend to enlist Western scientists with appropriate experience. NEC and Fujitsu have established facilities as well as collaborative research with groups in the United States, Australia, and Europe for this purpose. Thus Japanese vendors are becoming more effective at accessing expertise beyond their shores.

For those users who need to write programs, standard languages run on all Japanese supercomputers, and the vendors scrupulously meet all announced standards while offering various enhancements and extensions. Optimized vendor libraries are also obtainable, with routines for such tasks as matrix manipulation and fast Fourier transforms. Japanese companies have large teams of programmers developing these libraries, and they also support well-known commercial libraries from the West, not to mention noncommercial projects such as Eispack and Linpack.

If the user interfaces are standardized, then portability is maintained along with efficiency. But nothing is being done in Japan with an eye toward standardizing scientific software. And almost no research comparable to that in the West is under way on portable numerical algorithms, as typified, for example, by the Lapack project at the University of Tennessee and other cooperating institutions.

Nor is there much demand for standardized software in Japan; vendors and users still develop libraries and user interfaces for their own platforms and applications. Japanese computer users write their own application software, and people who



have studied it from the inside claim it is often quite good. In summary, Japanese software for supercomputers, especially for multiprocessors, is only now emerging.

#### **PARALLELISM: EARLY START, SLOW PROGRESS.**

Research into parallel computing started in Japan about 15 years ago at universities and national laboratories. As in the United States, early interest in the technology was tied up with physics, and this branch of parallel computing spawns new ideas and architectures to the present day. Special-purpose machines have been built for the study of quantum chromodynamics, gravitation, and thermonuclear plasmas, and other subjects.

The first noteworthy parallel computer developed in Japan was probably PAX. Built at the University of Tsukuba in Ibaraki-ken in the late 1970s, it still stands as one of Tsukuba's most important and influential achievements. It was designed for work in nuclear engineering, but succeeding generations of this architecture have been directed toward more purely scientific uses. The latest version is a specialized machine for quantum chromodynamics applications called QCDPAX. It has achieved 8-13 gigaflops with 288-480 nodes. An operating system is lacking, and the only language is a parallel version of C.

Many other parallel computers serve such special purposes as neural network simulation, digital signal processing, checking the logic of planned semiconductor circuits, fingerprint matching, deoxyribonucleic acid (DNA) sequence search processing, and image or graphics processing. At least 50 such special-purpose systems have been developed. And at least 50 very varied research projects on parallel computing are under way, mainly at universities.

Japan's many special-purpose parallel designs have engendered a large overlap between chip and prototype parallel-computer development. Nonetheless, these computers are mostly stepping stones for further work and will never be produced commercially. Indeed, no Japanese manufacturer currently offers a commercial massively parallel computer, although almost all of them can boast one or more experimental models and one or more of these models could be marketed in the near future.

Several obstacles exist to faster progress in the development of commercial massively parallel architectures in Japan. One is the general feeling that conventional (modestly parallel) systems have advantages and further potential. Others include a limited market for conventional supercomputers, recent evidence of downsizing of computer systems, and predictions that workstations will flourish most in the computer market.

Another, more profound, obstacle is the lack of small venture start-ups in the country's computer industry. But the large vendors are loath to develop the architectures out of fear for the markets for their existing products. Perhaps the developmental role

will be played by other manufacturers who have till now been preoccupied with home electronics; some of these are eager to move into the high-performance computing business and are slowly succeeding.

Nonetheless, the failure of Japanese parallel computers to reach the West must not be construed as evidence that the projects are not active or well established. For the most part, a lack of resources has dictated that most of these systems satisfy educational needs, rather than blossom into commercial products. Given this poor hardware availability, plus the well-known Japanese interest in building systems, it is only to be expected that the kind of research into parallel applications so common in the West is much less common in Japan.

A complete list of Japanese projects in massively parallel architectures is impossible here, but descriptions of a few of the more important will give a feel for the scope of Japanese commitment and progress in this area. As most machines are still under development, details of their specifications and performance could only be obtained from the

developers and may be a bit overstated.

NEC is working on two massively parallel computers, HAL and Cenju. HAL is a logic circuit simulator, with 31 special-purpose processors. Circuit models can be described for it in a high-level language, making simulation easy. NEC uses HAL for routine logic and fault simulations and claims that it is more than 10 000 times faster than conventional gate-level software simulators. At present HAL III, the third version of a multiprocessor electronic-circuit simulator, is in operation.

Cenju (which derives its name from the 1000-handed "sen-ju" goddess Kannon) is a computer with multiple instruction, multiple data (MIMD) streams. It was originally designed for the simulation of circuit transients using the Simulation Program with Integrated Circuit Emphasis—that is, generating and solving the systems of ordinary differential equations associated with voltages and currents in semiconductor circuits. Cenju uses a combination of distributed and shared memory.

The many processors sometimes commu-

### **How MITI pushed parallel machines**

As with other critical technologies in Japan, the development of parallel computers there was greatly assisted by Government-sponsored consortia—in particular, two national projects promoted by Tokyo's Ministry of International Trade and Industry (MITI) in the 1980s.

One was the nine-year Superspeed project, begun in 1981 with a goal of producing a scientific computer capable of 10 billion floating-point operations per second (gigaflops). There were two main subprojects, one for the development of high-speed novel devices and the other for architectures. The project concluded in 1990 with the demonstration of a Parallel Hierarchical Intelligent computer, which had four processors and did in fact achieve a peak processing rate over 10 gigaflops and a real performance over 1 gigaflops.

The prototype could not be directly commercialized, and the results have not yet been incorporated into commercial projects. Nonetheless, Superspeed was considered helpful by the various Japanese companies that participated.

The other national project, which was more focused on parallel computing, was the central research Institute of New Computer Technology (ICOT), which was founded in 1982 with Kazuhiro Fuchi as scientific director. He and four group leaders from the Electrotechnical Laboratory and NTT Transmission Systems Laboratory, together with two dozen young researchers from industry, began to work in a field unknown to Japan (eventually about 100 researchers were working at ICOT). To conclude the project, five different parallel inference machines were built by the cooperating Japanese companies of Fujitsu, Hitachi, Mitsubishi, Toshiba, and Oki Electric Industry Co., also based in Tokyo. The largest had 512 processors and achieved about 100 megalips (millions of logical inferences per second).

Although hardware was extensively developed, any

lasting impact will result from software. ICOT wrote a collection of about 70 sample parallel applications, including some for the routing of large-scale integrated (LSI) chips, protein sequence analysis, and legal reasoning; all are examples of symbol processing with low data parallelism—quite unlike most supercomputing applications.

At the final ICOT conference last June 1-5 in Tokyo, MITI announced that it would make all ICOT software available to the general public in Japan and abroad without charging the usual license fees. For the moment, the software will run only on special hardware, but preparations are being made for its migration to other environments. This is an intelligent step by MITI, as otherwise the intellectual products of the fifth-generation computing systems (FGCS) would be likely to remain unused outside Japan. Probably this new policy will also be applied to the Real World Computing project, MITI's new 10-year project [see p. 47].

ICOT failed in its goal of launching fifth-generation computing systems. Traditional sequential and vector architectures have proven to be more robust than was thought. Users and industry are insisting on systems compatible with existing resources (hardware makers do not plan to commercialize it). Nevertheless, ICOT's logic programming systems are probably the best in the world.

One of the lasting results of the project was the training of many young Japanese in symbolic computation, software, and parallel computing; there is probably as much or more expertise within Japanese companies on these topics as within any Western counterparts. FGCS provided both money and focus to successfully lubricate the start-up of a knowledge-processing industry in Japan. The basic research from Japan in theorem proving and related areas is now comparable to that in the West, a big change from 10 years ago. —D.K.K. and U.W.



nicate with one another in Spice-type problems, but a common bus with a high data-transfer rate suitable for use by many processors was difficult to build. Instead, Cenju is designed in a tree-like form of eight clusters of eight processors. Clusters are connected by a multi-stage network. Each processing node has a 20-MHz MC68020 Motorola CPU with 4M bytes of dual-ported local memory and a fast Weitek WTL1167 chip for floating-point arithmetic. The design allows up to 256 processors.

Cenju met its goals, but the 68020s limited its performance. The architecture traded off simple hardware against complex system software. Cenju uses a special operating system, C or Fortran, and a parallelizing compiler.

A second-generation machine, Cenju II, is now in gestation. It has two MIPS R3000 microprocessors in each processing node, one each for application execution and communication control. Its target specifications are 6.4 gigaflops using 256 nodes, with a total of 16G bytes of memory. A 16-node prototype has been completed. With a full-scale Cenju II, NEC could enter the market, possibly joining near-term offerings from Fujitsu and Matsushita Electric Industrial Co., based in Osaka.

At Fujitsu, a development effort begun in 1987 culminated a little over a year ago in the AP1000, a single-user MIMD computer. Its 64-1024 processing nodes are arranged in a two-dimensional torus. Intended for applications in general science, the AP1000 is connected to a host, currently a Sun 4/390 workstation. To compute on the AP1000, a user must write two (or more) programs, for the host and the nodes (in C or Fortran). Communication is by calls to library routines. A subset of Unix runs on the computer's nodes.

Communication between the host and the AP nodes is by way of a 3-megabyte/s host interface (VMEbus), into a 32M-byte buffer. This link is much too slow for rapid graphics, something Fujitsu has promised to rectify within the next year.

Each cell is composed of 16M bytes of dynamic RAM, a 128K-byte cache (equivalent to 16K 64-bit words), a 25-MHz Sparc chip designed by Sun, a Weitek floating-point unit, assorted controllers, and network interfaces. Each cell has a peak performance based on the Weitek chip of about 8.33 megaflops (32-bit) or 5.56 megaflops (64-bit).

Fujitsu is collaborating with the Australian National University (ANU) in Canberra, where one AP1000 has been installed since summer 1991. ANU staff are engaged in a variety of research and report on this at a yearly Fujitsu/ANU Cellular Array Processor Symposium. A similar system is being installed at a university in the United Kingdom. Richard Brent of the university's computer science department has already obtained almost 300 megaflops on a 64-node AP solving linear equations in double pre-

cision. Peak for a 64-node machine is about 350 megaflops.

In an unexpected disclosure last July at the Sixth International Conference on Supercomputing in Washington, D.C., Fujitsu revealed that it was working on a massively parallel system capable of supporting up to "a few hundred" processors and a peak processing rate in excess of 300 gigaflops. The machine's architecture will combine vector processing with parallelism, Fujitsu said. Each node of the system will include a vector and a scalar processor and a memory unit "configured in a shared architecture," according to the company.

Fujitsu executives in the United States said the machine, whose name was not disclosed, would be marketed in that country. The schedule for its commercial release was not made known, however.

Possibly closer to the market is a system from Matsushita, the consumer electronics giant that also represents Cray Research and the Aachen-based German firm Parsytec GmbH in Japan. Adenart was developed in conjunction with Tatsuo Nogi from Kyoto University. Its 256 processors form a 16-by-16 array and are connected by a three-dimensional network designed for alternating-direction implicit (ADI) methods, a numerical technique useful in solving discretized partial differential equations. This network is also said to be effective for general-

purpose data transmission, so the application fields go beyond three-dimensional partial differential equations. Adenart's peak performance is 2.56 gigaflops and Nogi estimates that its effective performance is 1 gigaflops. As a high-level language, a parallel expanded Fortran called Adetran has been developed.

Adenart is very compact: it measures (excluding its host) 50 by 139 by 70 cm and should have a good cost/performance ratio. For the present, the machine is open for experimental use by universities and other public organizations.

Future plans go in two directions: to build a smaller Adenart-64, a fast workstation version, and to upgrade the system. In 1992, the first OHM series model will be announced. The OHM256 will have 256 faster processors with a 25-gigaflops peak, and could sell for around 300 million yen. The firm will probably connect four OHM256s and develop a 100-gigaflops machine, the OHM1024.

In 1991 Sharp published an outline of its data-driven processors (DDP), claiming that this type outperforms complex- or reduced-instruction-set computers (CISC and RISC, respectively). This document is a clear and well-written statement of the company's plans. The present DDP was jointly developed with Mitsubishi. In this machine, 1024 processing nodes can be connected for applications in image processing,

## Data flow is 'in' in Japan

Logic programming and data flow are two areas in which Japanese parallel processing is comparable to, or even ahead of, that in the West.

Data flow is based on the simple idea that an instruction should be executed when its data become available. Machines that use a single instruction such as + or \* as a data flow node are called pure data flow. Some examples in Japan are the ETL Sigma-1, the NEC ImPP, and the recent Sanyo Cyberflow.

However, pure data-flow machines are somewhat inefficient, so control flow has been introduced, although this requires extra efforts in compiling. Future data-flow computers will be hybrids, combining the best aspects of data-flow and efficient execution of sequential code.

Several groups in the West are also working on data flow (the Laboratory for Computer Science at the Massachusetts Institute of Technology [MIT] in Cambridge is one of the best known), but the Japanese lead in building hardware; they are behind in data-flow software. Western researchers disagree about the practicality of data flow, or how much real progress has been made, especially in the area of software. However, for the data-flow approach to succeed, a great deal of work in languages and compiling technology will be necessary.

The Electrotechnical Laboratory (ETL) in Tsukuba, Japan, has a long data-flow tradition. Besides the finished SIGMA-1 project discussed in Section 3, there is the EM-series. EM-4 development started in 1986, following the EM-3 model (1981-

85). In line with the older machines, a modified data-flow control scheme was applied. By introducing various improvements, a 1000-processor-element machine might be feasible, but probably will not be built as part of this project, which ends in 1993.

In a recent paper, Andrew Shaw from MIT compiled a subset of \*Lisp into C and then produced object code with the EM-4 compiler. His performance results for an (integer) radix sort are superior to the performance on a 2048-node CM-200, from Thinking Machines Corp., Cambridge, Mass.

The ETL group is studying the development of the next model, EM-5, with higher integration and much higher performance. They are trying to develop it as a general-purpose machine with 16 348 processing elements, using a new 80-million-instruction-per-second processor, EMC-G. EM-5 will depart further from pure data flow. Using EM-5 a teraflops machine might be feasible, and will probably be built for the Real World Computing project.

Another ardent promoter of the data-flow principle is Hiroaki Terada, a member of the electrical engineering department at Osaka University. He began developing the Q series (Q stands for queue) in the '80s. Q-x employs the data-flow principle, realized by a ring-type machine. After initial governmental support, Terada convinced Sharp Ltd. and Mitsubishi Electric Corp. to step in. They developed the basic ideas together but are now working separately. Sharp recently announced a one-board data-flow processor based on this concept.

-D.K.K. and U.W.



graphics, and so on. A programming language similar to C is compiled so that parallelism is extracted automatically. Software development is supported by the usual tools.

Sanyo Electric Co. of Moriguchi City in Osaka Prefecture is another leader in consumer electronics that did not want to be left out. Its Enhanced Data-Driven Engine (Edden) has up to 1024 processing nodes linked in a two-dimensional torus. Each of the nodes is based on a 32-bit, custom-designed one-chip CMOS processor. In early 1992 Sanyo announced a commercial version, called Cyberflow. It is sold as a small desktop system with 4-64 nodes, with a claimed peak of up to 640 megaflops. A parallel assembler and a C compiler are provided. Applications are image processing, graphics, and simulation.

Hitachi and Mitsubishi cooperated with other companies on the construction of different experimental parallel machines while participating in Government-organized consortia, but neither has its own commercially oriented, indigenous, massively parallel development effort at present. However, Hitachi is planning to build a 300-gigaflops machine for the newly opened Center for Computational Physics at Tsukuba University. The design, not yet made public, may build upon the extensive PAX experience within Tsukuba and Tokyo universities.

At Toshiba Corp., headquartered in Tokyo, Shigeru Oyanagi developed a system, called Prodigy, for symbol processing and computer-aided design. It is a 256-node machine featuring an eight-cube where each node can communicate with only three neighbors. This arrangement reduces the performance somewhat, but allows inexpensive construction. And at the Tokyo-based NTT Transmission Systems Laboratory, Sadayasu Ono is developing a massively parallel system for signal processing called NOVI.

Despite these and many other development efforts, there are only a handful of operating parallel computers in the hands of industrial software developers, even fewer at universities, and very few available to users who want to solve real problems. This situation seems poised for change, as Fujitsu's AP1000 system is now making its way into the hands of Japanese researchers.

Research in software for parallel computing is far behind that in the West but is growing, and will accelerate further as more systems get into researchers' hands. There is also significant work in distributed computing, which can be performed on networks or clusters of workstations. The large companies are engaged with (vector) supercomputers. Lacking parallel applications, they are cautious about moving, commercially, into even more risky massively parallel systems.

Other researchers, such as Akinori Yonezawa at Tokyo University, are concerned with software for massively parallel

systems in general, experimenting with different existing Japanese machines. This cross-evaluation marks a necessary departure from isolated, self-contained developments, and is a first step toward real applications.

**GETTING REAL.** A cornerstone of future research in parallel computing is the ambitious Real-World Computing Program (RWC), which has aroused tremendous interest in the West. This is to be a 10-year Ministry of International Trade and Industry (MITI) program with a budget of about 60 billion yen (almost US \$500 million), focusing on R&D

## No Japanese manufacturer now offers a commercial massively parallel processor, but that will soon change

for what is called flexible or intuitive information processing—the way human beings absorb information and make decisions.

The program organizers envision the need for a substantial computational base comprising both massively parallel computing systems and neural systems. These might be physically distinct, although it is hoped that they will be integrated into useful platforms. Research into optical computing is planned as part of the technology base for both systems, either for interconnection or if possible for actual processing. Thus research will be supported on all aspects of these systems, but the computational base is seen mostly as a platform upon which more theoretical and functional work will be done to support the higher-level research thrusts of the program.

Work on the massively parallel system will involve all aspects. However, scientists working on data-flow computer architectures are to be key personnel on this part of the project, so it is expected that their ideas will play an important role [“Data flow is ‘in’ in Japan,” left]. During the first of the project's two phases, various competing systems will be developed with as many as 10 000 processors. During the second phase, the best ideas would be retained and eventually a million-processor system built. The EM-5 already being implemented at the Electrotechnical Laboratory in Niihari-gun, Ibaraki-ken, affiliated with the MITI agency for Industrial Science and Technology, is planned to have 64 000 processors, hopefully achieving teraflops speed.

Research in neural systems will focus on software and hardware, in two phases, as above. The final goal is to develop a one-million-neuron network composed of 1000 subnetworks of 1000 neurons each, with a

total processing speed of 10 teracups (10 trillion connection updates per second).

A novel feature of the real-world computing project is that it is to be international, interdisciplinary, and open. A research association, called RWC Partnership, has been set up by a half-dozen Japanese computer manufacturers, but foreigners may participate. To aid this effort, Japanese intellectual property rights laws have been altered, and project documents as much as possible are to be issued in English. Collaboration between the United States and Japan is not yet established, but probably will be within the subarea of optoelectronics.

With such projects as RWC, Japan will strengthen and diversify its position in parallel processing. At the same time, semi-commercial massively parallel systems will become more widely available to Japanese researchers, accelerating progress and paving the way for the introduction of commercial systems. It can safely be assumed that Japan, with its traditional inclination to autarky, will stay in the race for teraflops machines. Quite possibly, it will dazzle the world once again with its

prowess in a technology critical to economic progress in the coming years and beyond.


**ABOUT THE AUTHORS.** David K. Kahaner has been at the U.S. Office of Naval Research's Tokyo office since 1989. He is currently on leave of absence from the Center for Computing and Applied Mathematics at the National Institute of Standards and Technology, formerly the National Bureau of Standards, where he was a group leader from 1978 until 1989. From 1968 until 1979 he was in the Computing Division at Los Alamos National Laboratory, New Mexico. His interests are numerical analysis and scientific computation, on which he has published many papers. In Tokyo he is analyzing Japanese and Asian high-performance computing activities, and distributing reports by electronic mail. He holds a Ph.D. in applied mathematics from Stevens Institute of Technology in Hoboken, N.J.

Ulrich Wattenberg is head of the Tokyo office of Gesellschaft für Mathematik und Datenverarbeitung (GMD) MBH, the German National Center for Computer Science. Together with the European Research Consortium of Informatics and Mathematics (Ercim), GMD acts as a think tank for the European computer industry. In graduate school at Philipps University, Marburg, Wattenberg's research interests were in crystallography, and his Ph.D. thesis was on lattice distortions in pure silicon. While working on the thesis, he stayed for three months at IBM's East Fishkill, N.Y., Laboratory. He also did research and studied the Japanese language at Tokyo University from 1974 to 1976, before moving to the quasi-governmental German Institute of Documentation, which is now part of the GMD. He established the Tokyo Liaison Office, which he is still heading, in 1977. ♦



# Balancing resources

*Gigaflops alone do not a system make; both massive storage and fast communications are also essential*



Discussions about high-performance computer systems usually emphasize the raw, peak processing power of the supercomputer—and understandably so. After all, the machine itself is the central tool, the most visible element, in a high-performance computing environment.

But it is not alone. Other, less visible, components in the environment limit the supercomputer's performance and ultimately determine how useful it can be.

Ideally, a computer's capability is defined by the size of its memory in combination with its computational processing rate. The memory size determines the complexity of applications that can be solved (what will fit into memory) and the processing speed determines how long the solution will take.

But reality is much more complex. At the National Center for Supercomputing Applications (NCSA) at the University of Illinois, solving a problem with a supercomputer is seen to involve more than just processing the numbers within its memory. In many cases—most scientific applications, for example—the desired result of a set of calculations is not raw data but information in a form that can be examined and understood by a human being.

Today, converting supercomputer output into that desired form generally means using several postprocessing computers as well as networks to move large quantities of data between the supercomputer, the postprocessors, and high-resolution display devices. In addition, storage systems are required to hold the data that is not being processed at a given time.

For a computing system to work optimally, the capabilities of all its major components must be in balance. Unfortunately, the balance between supercomputer capacity

(processing rate and main memory size) on the one hand and local-area network and storage system capacities on the other is questionable and getting worse. Capabilities in data communications and mass storage systems have simply not kept up with those of supercomputers themselves, so much so that the former are poised to become serious bottlenecks in the not-too-distant future.

**KEY TRENDS.** To amplify, in the 12 years since 1980, peak potential processing power and the memory sizes of supercomputers have increased by three orders of magnitude [Fig. 1]. Yet, over the same period, local-area network (LAN) communication rates have increased by less than two orders of magnitude.

A bit of care is needed in interpreting the graphs of Fig. 1. In particular, note the two data points for processing rate in 1992. The lower one (16 billion floating-point operations per second, or gigaflops) is for a Y-MP C90 vector multiprocessor from Cray Research Inc., Eagan, Minn. The upper one (131 gigaflops) is for a 1024-node CM-5 from Thinking Machines Corp., Cambridge, Mass. To approach peak potential processor throughput on either machine, the structure of the application must be optimized to take advantage of the machine's vector facilities and multiple processors. The actual processing rate that an application achieves on these machines depends on how well the application maps to the machine's architecture.

**With raw computing power on the rise, system performance is bumping into network and storage limits**

Another area, too, has contributed to the demand for faster communications. While LAN communication rates were falling behind memory size and processing speeds in the 1980s, the metacomputer emerged. That concept refers to the use of several computers working together simultaneously on a single application. It has come into favor as users have recognized that the solution to a supercomputing problem is generally bound up with many tasks, administrative

and computational, both before and after the supercomputer run.

A metacomputer can be as simple as a single supercomputer connected to a single workstation in a client-server relationship. Or it may be complex enough to embrace several supercomputers exchanging large arrays of data as they cooperate on a numerical problem.

As for storage, experience at NCSA has shown that, at least for advanced scientific applications, output files tend to expand in direct proportion to main memory growth. Since mass storage systems save those files, requirements for mass storage track main memory size, and are already expanding well beyond what traditional disk farms can feasibly handle.

Clearly, data networks and storage systems must be improved to keep pace with supercomputer developments. Networks capable of transmitting tens of gigabits of data per second must be developed to replace megabit-per-second networks, and petabyte archives will be needed to replace those capable of storing mere terabytes.

Work on those objectives is progressing. The required network technologies and architectures are being developed in high-performance computing centers as well as in wide-area network testbeds. New storage technologies are also emerging along with changes in the way that large data archive systems are designed.

**NETWORK NEEDS.** Just how fast should a local-area interconnect network be in order to support meta-applications running on high-performance computers? One answer considers the peak potential output that a particular application or meta-application component might generate on a given supercomputer. That peak could be reached if, for example, all the data were to be sent to another supercomputer in real time for post- or co-processing, or if it all were to be sent to a storage device in real time.

To determine how much data (new information) is produced by a given application, consider numerical simulations of physical systems. In this common type of application, a physical system like an aircraft wing or a weather pattern is represented by an array of zones stored in memory, each representing an area or volume of the physical system. Each zone consists of a set of variables (floating-point numbers) that represent the physical state of that zone at a single point

Charles E. Catlett  
National Center for Supercomputing Applications




# IF YOU THINK YOU'RE GOING TO GIVE US A PROBLEM-

We've been turning trouble into opportunities for nearly half a century. We don't get mad - we get technical. That's how the Sarnoff Research Center came up with the first complementary metal oxide semi-conducting chip. And the first solid state amplifier for space use . . . . The problem is, you should be up front with us - so we can use our heads, hunches, and tremendous reservoir of knowledge and experience in such areas as superconductivity, surface emitting diode lasers, and high-definition TV systems . . . . We'll give as good as we get. For industry and government, we use the genius of the mind on basic research and product-oriented, market-driven problem solving. Heads in the clouds, feet on the ground, we're ready for the toughest challenge you can throw at us . . . . For information contact: Public Affairs, David Sarnoff Research Center, CN 5300, Princeton, NJ 08543-5300 (609) 734-2507. So what's your problem???

## TERRIFIC!

DAVID  
**Sarnoff**  
RESEARCH CENTER  
Subsidiary of SRI International


FIFTY YEARS OF VISION




Circle No. 34



Frequency responses of a family of control systems, modeled, simulated and visualized in MATLAB 4.0.



# We see your expectations of visualization and we raise them.



Three views of the surface height of a penny show user customizable object-oriented graphics in MATLAB 4.0. Data courtesy of NIST.

Combine advanced visualization with the powerful computation of MATLAB, and gain new insight into your most challenging problems.

## **The MathWorks introduces MATLAB 4.0**

MATLAB 4.0 blends visualization techniques and numeric computation into a seamless interactive environment that redefines how you can solve complex problems. You can analyze data numerically and visually, simulate models and see the results immediately, or explore ideas and test them interactively.

### **More than meets the eye**

MATLAB 4.0 provides you with an extensive library of built-in computational tools, combined with a powerful fourth-generation language.

It thereby offers the convenience of a pre-packaged application program and the extensibility and flexibility of a high-level language. Much easier to use than Fortran or C, MATLAB 4.0 yields tremendous gains in productivity and creativity.

New high-level functions enhance these gains—functions in MATLAB ranging from advanced graphics to powerful sparse matrix support to integrated program development tools.

Add SIMULINK to MATLAB 4.0 and you can perform dynamic system simulation of nonlinear models in a graphical mouse-driven environment. And with specialized application toolboxes, you can tailor the MATLAB environment to address your specific needs.

We've devoted the next page to giving you a more complete picture.



# MATLAB® 4.0

## Picture the Power

### Broadly applicable, widely used


Over 100,000 technical professionals use MATLAB in diverse applications areas, including signal processing, control engineering, chemistry, economic modeling, and data analysis. They use it as their primary analysis tool in industrial, academic, and government settings in over 50 countries.

### Flexible & extensible

With the open system architecture of MATLAB, you can see the algorithms, edit them, and create your own to address specific needs. MATLAB provides multi-platform interoperability, so you can use your applications and data across platforms without modification. As a result, MATLAB has become the ideal way to create and communicate algorithms and ideas—a *lingua franca* for technical professionals.

### Easy to use

MATLAB has always been powerful, yet easy to use—a fast and accurate number-cruncher, a high-level prototyping environment, and a complete programming language. MATLAB 4.0 adds many new features that make it more powerful and even easier to use.




### New feature highlights

The most visible changes to MATLAB are graphical.

**NEW**

With our new object-oriented Handle Graphics™ system, you can customize virtually every conceivable attribute of your MATLAB plots.



Riemann surface of the complex cube root function shows the capability of MATLAB 4.0 for mathematical visualization

New graphics capabilities include:

- 3-D shaded color surface graphs
- 3-D contour plots
- 3-D data trajectories
- Image display
- Light sources
- Faceted and interpolated surface rendering, plus texture mapping
- Animation
- Color and black & white PostScript® output (Level 1 & Level 2)

Image representation of Southeastern New England altitude data created in MATLAB 4.0. Data courtesy of NOAA.

### Beyond the visual

Other new features in MATLAB 4.0 include:

- **Over 100 new functions**
- **Graphical user interface controls**  
Add "widgets"—such as sliders, buttons, and menus—to create customized MATLAB graphical user interfaces easily.
- **Enhanced file I/O**  
Import and export your data in any format.
- **Sound output**  
Play data and functions as sound (SPARC only).
- **Integrated debugging environment**  
Create MATLAB functions and scripts more quickly and easily.



Spectrogram of Handel's Hallelujah Chorus, computed and displayed with MATLAB 4.0 and the Signal Processing Toolbox.

### Powerful, versatile toolboxes

Application toolboxes, designed and written in MATLAB by world-class experts in their fields, provide specialized solutions. They're based on MATLAB, and they combine pre-packaged functionality with an open systems approach that allows you to see the algorithms and to modify them to suit your needs.

The comprehensive MATLAB toolbox family includes:

- Signal Processing
- Control System Design
- Robust-Control Design
- Mu-Analysis and Synthesis Design
- System Identification
- Neural Networks **NEW**
- Nonlinear Optimization
- Spline Analysis
- Chemometric Analysis

MATLAB 4.0 is available now for Sun SPARCstations, and will soon ship for other standard MATLAB platforms including: 386/486 PCs, Macintosh, HP/Apollo, DECstation, VAX/VMS, IBM, Silicon Graphics, Convex, & Cray. **For more information, please call us at (508) 653-1415.**

**The MATH WORKS Inc.**

24 Prime Park Way/Natick, MA 01760  
Tel: 508/653-1415 Fax: 508/653-2997  
Email [info@mathworks.com](mailto:info@mathworks.com)


France: Scientific Software, +33-145-34-23-91  
Germany: Bausch-Gall, +49-89-323-2625  
Israel: Omikron Delta, +972-3-561-5151  
Japan: Cybernet Systems, +81-3-3982-4641  
Scandinavia: Comsol Europe, +46-8-15-30-22  
Switzerland: Comsol AG, +41-31-961-70-11  
U.K.: Cambridge Control, +44-223-420-722  
U.K.: Rapid Data Ltd., +44-903-202-819



# Transactions from Japan covering

Electronics, Communications, Computer Sciences

Published by The Institute of Electronics,  
Information and Communication Engineers (IEICE)



## Facts about IEICE:

Largest Japanese professional society in engineering with 39,000 members.

Established in 1917.

12 monthly publications; 1 Journal, 7 Japanese Transactions, and 4 English Transactions.

Sponsors/Co-sponsors many international conferences; eg. ISS'92, Networks '92, OEC'92, Globecom'87, ICASSP '86, ISCAS'85.....

## IEICE Transactions are the English publications of IEICE,

featuring Original scientific and technical papers  
Tutorials and invited papers  
Special Issues/Sections on hot topics

## Over 5000 pages of valuable information printed in 4 volumes.

- Vol.A: IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences
- Vol.B: IEICE Transactions on Communications
- Vol.C: IEICE Transactions on Electronics
- Vol.D: IEICE Transactions on Information and Systems

## For subscription or for further information, contact Maruzen

Maruzen Co., International Div., Export Dept., P.O. Box 5050 Tokyo International 100-31 Japan

## Annual Subscription Fee

For each volume (12 issues) \$100 / For all 4 volumes (48 issues) \$360



in time. To describe the evolution of a system over time, a supercomputer solves the equations that govern the system in an iterative fashion, updating the variables in each zone at successive points in time (timesteps).

The application creates a new array of zones and variables each time the equations are solved for the entire system. Each application involves a unique set of equations and numerical algorithms, which together determine how many floating-point operations are required, on average, to update a single variable in the simulation, thereby creating new information. By dividing the rate at which the application actually runs on a given supercomputer (the number of floating-point operations per second) by the number of operations required to update a variable, we can find the average number of variables updated per second. The peak potential I/O rate is then found by multiplying that average by the number of bytes or bits per variable.

Note that a maximum I/O requirement calculated in this manner assumes that all created data is also transmitted. The number is essentially related to the rate at which a given supercomputer works its way through a problem updating variables. It is unrelated to the problem size, which will combine with this rate to yield the total solution time for the complete evolution of the physical system.

To better see how the various parameters are interrelated, an analysis was made of a particular problem—an atmospheric simulation model, to be specific. The results of that analysis showed that, on average, variables were updated every 200 floating-point operations. Using four Cray-2 processors simultaneously on this application in 1989, the actual processor throughput achieved was roughly 400 million floating-point operations per second (megaflops). Since each

## Defining terms

**Datagram:** a packet or cell.


**Fibre Channel:** a fiber-based standard, under development in ANSI X3T9.3, for high-speed local-area networks and metropolitan-area networks with rates varying from about 100 to 1000 Mb/s.

**Hippi:** a high-performance parallel interface, a copper-based data-communications standard developed by Task Group X3T9.3 of the American National Standards Institute (ANSI), capable of transferring data at 800 Mb/s (32 parallel lines) or 1.6 Gb/s (64 parallel lines).

**Potabyte:**  $10^{15}$  bytes.

**Terabyte:**  $10^{12}$  bytes.

**Virtual circuit:** a circuit-like service provided over packet-switched networks, which makes use of sequence information and retransmission or error correction to ensure that data arrives reliably and in the proper sequence.



[1] Growth in local-area network capability has fallen behind advances in processing power and memory capacity. The parameters shown here are normalized to 1980 figures of 100 million floating-point operations per second, 8M bytes, and 50 Mb/s. The processor rates are peak potential (aggregate, not necessarily observed) figures for available Cray supercomputers. For 1992, two figures are given: 131 billion floating-point operations per second for a 1024-node Thinking Machines CM-5 and 16 billion for a 16-processor Cray Y-MP C90 vector multiprocessor. Note the potential difference in peak rates between massively parallel and vector supercomputers.

variable was represented as a 64-bit floating-point number, the peak potential I/O rate was 16 megabytes/s, or 128 Mb/s [Fig. 2].

To run this same application on the CM-5 of Fig. 1, assuming that the algorithm complexity remains constant (200 operations per variable), and allowing for an actual processor throughput of 10 percent of peak (13 gigaflops), the peak potential I/O rate of this application becomes 4.16 Gb/s.

More typically, a user might keep just enough data to analyze the output, maybe 10 percent of the data—say, every 10th timestep—at half precision (32 bits). Under such circumstances, the required bandwidth would be reduced to 5 percent of the peak, or 208 Mb/s for the CM-5 example above. Examples of this type of work abound today. **SUFFICIENT STORAGE.** Tough as they are, the network requirements of these state-of-the-art supercomputers are no steeper than their storage requirements. Many advanced scientific applications generate sets of tens to hundreds of output files, any of which might occupy 25 percent of available memory. (The 25 percent figure arises mainly because numerical simulations keep three copies of their variable arrays in memory: the previous timestep, the current timestep, and the timestep being computed.)

As recently as two years ago, the largest supercomputer memory available was the Cray-2's 4G bytes. With this size, a large application might create 100 files 1G byte long for a total of 100G bytes of data. With some difficulty, that data might be analyzed on a scientific workstation with plenty of disk

space, perhaps downloading one or two of the files at a time, depending on how much workstation disk was free.

Consider, though, a supercomputer with 32G bytes of memory—the CM-5 with 1024 nodes, for example, or the Cray M90. The same application, scaled to available memory, will now produce 100 files, each 8G bytes long, for a total of 800G bytes of data. At this point, the rest of the computing environment is clearly out of balance.

Today's supercomputers might have 50–100G bytes of disk capacity; a high-performance mass storage system might have 200G bytes. The 800G-byte data set, then, has outgrown rotating disk, not to mention a workstation. To run this application, therefore, data must be staged to the supercomputer disk and maybe also to a central archive disk. In the end, though, it will have to be held on a cheaper, secondary storage medium because not enough disk space is available locally to hold all 800G bytes.

Unless there is enough disk staging area to store the entire data set, or unless the problem is to be broken into several small runs (with data shuffling to and from secondary stor-


age in between), the secondary storage will need to handle data at a very high rate—at worst, the peak potential rate of Fig. 2, but more typically, 5 percent of that rate, as noted earlier.

**NETWORK NEWS.** Ten years ago Cray supercomputers were generally connected to front-end processors via HYPERchannel, a proprietary network with a peak throughput of approximately 50 Mb/s but peak individual transfer rates of 5–10 Mb/s. Today, the fiber distributed data interface (FDDI) is replacing this network in many supercomputer centers, and this trend is expected to continue.

FDDI host interfaces on supercomputers and workstations currently operate in the 20–50-Mb/s range. As with Ethernet in the 1980s, however, enhancements to FDDI host interface implementations will doubtless bring individual transfer rates up to a much higher fraction of the technology's 100-Mb/s aggregate capacity over the next several years.

Within the last decade, interface improvements have resulted in several new products. In 1987 demand for higher I/O rates from Cray supercomputer customers led to the development of the HSX channel interface, an 856-Mb/s parallel channel. At the same time the high-performance parallel interface (Hippi) standard began to take shape within ANSI's Task Group X3T9.3 under the leadership of the Los Alamos National Laboratory in New Mexico. The Hippi physical layer standard (Hippi-PH) defines 32- and 64-bit parallel interfaces that run at rates of





[2] The maximum I/O rate is a function of algorithm complexity and processor throughput. The examples shown are for an application that averages 200 floating-point operations per variable calculation. When run at 400 megaflops, its peak potential I/O rate is 128 Mb/s; at 26 gigaflops, that figure rises to 8.3 Gb/s.

800 Mb/s and 1600 Mb/s, respectively.

In addition to the physical layer standard, the ANSI X3T9.3 group is nearing completion of further Hippi functionality documents [Fig. 3], with several in the review and final approval phase of standardization. These include framing protocol (Hippi-FP), which defines the packet header and packet structure; link encapsulation (Hippi-LE), which provides a mapping to IEEE 802.2 for support of common network protocols such as TCP/IP; and switch control (Hippi-SC), which defines the control and addressing of Hippi physical layer switches (usually crossbar switches).

Mappings are also being developed for the Intelligent Peripheral Interface (IPI-3) command sets for disk and tape, allowing support of technologies such as striped disks or tapes directly connected to a Hippi channel or LAN.

Whereas Hippi applies to copper cables and distances of 25 meters or less, a Serial-Hippi standard has been developed that will use optical fibers to extend Hippi to distances as long as 10 km.

**FLEXIBLE 'FIBRE.'** Another alternative fiber-optic technology, called Fibre Channel, is also taking shape, driven primarily by IBM Corp. and the Lawrence Livermore National Laboratory, Livermore, Calif. Fibre Channel builds on many of the concepts introduced by Hippi, and, like its progenitor, is being standardized within ANSI Task Group X3T9.3.

The differences between Hippi and Fibre Channel go beyond the differences between copper cabling and optical fibers. While the Hippi network approach focuses on a simple set of options, Fibre Channel seeks to supply a wide variety of services and capabilities.

For example, Hippi looks primarily at small groups of devices within a machine room, with 800- or 1600-Mb/s data rates. Fibre Channel, on the other hand, aims at switches with up to 4096 connections, distances of up to several kilometers, and multiple data rates (132.8, 265.6, 531.2, and 1062.4 Mb/s).

Moreover, the Fibre Channel activities encompass not only the lower-level network functionality similar to Hippi's, but also multiple types of connections (datagram and virtual circuit, for instance). While the Hippi physical layer interface standard is a 100-pin connector, the Fibre Channel physical layer has many options, including coaxial cable

and the choice of single- or multi-mode fiber, long- or short-wave lasers, and long- or short-wave light-emitting diodes.

Hippi development started about a year before the Fibre Channel efforts, and led to commercially available products within the year. In contrast, Fibre Channel products have been released only lately, and because of the ambitious and comprehensive nature

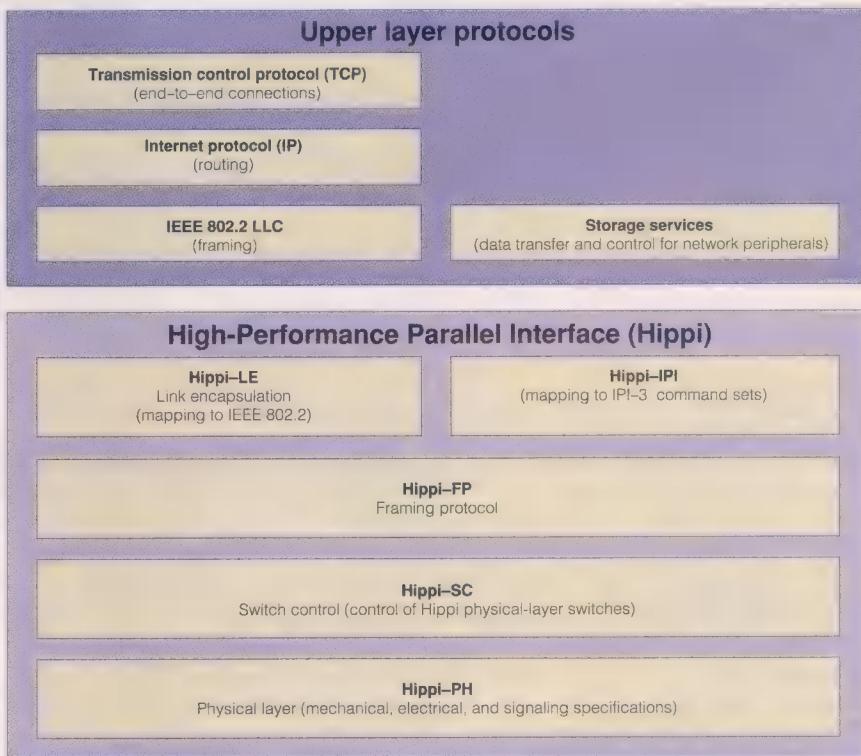
of the technology, even the physical layer document is at least a year from final approval as an ANSI standard.

Most recently, the Scalable Coherent Interface (SCI) standard, which has the potential for LAN speeds of 8 Gb/s, has been introduced.

**WANs LAG LANs.** Although LAN technology has not stayed abreast of advances in computer technology, it has made significant progress over the past decade or so. The same cannot be said for wide-area network (WAN) offerings, although this has begun to change recently. Over the past decade, the highest-capacity circuit available for lease has remained constant—the 45-Mb/s DS-3. Because of the expense of those circuits, most corporations today build data networks using 1.5-Mb/s circuits.

Most private networks in place today consist of a combination of leased circuits and customer-owned switching equipment (internet protocol routers, LAN bridges, multiplexers, and so on). The bursty nature of high-performance computing applications, however, requires a much more flexible structure for bandwidth management if it is to be at all cost effective.

Such a structure may be provided by a pair of technologies upon which the telecommunications industry is converging to build the infrastructure for its next generation of services: synchronous optical network (Sonet) and asynchronous transfer mode (ATM). Sonet is a hierarchy of trunk bandwidth rates ranging from 55 Mb/s up through 4.8 Gb/s, and will eventually go even



[3] With the Hippi physical layer standard approved, work is proceeding on the upper layers, including switch control, which defines the control and addressing of Hippi physical layer switches, and framing protocol, which defines the packet header and structure.



higher. ATM is a switching scheme involving fixed-length packets called "cells." In the telecommunications world of the next generation, Sonet transmission facilities will carry ATM cells between ATM switches. The cells will carry all types of traffic—voice, video, and data. The short, fixed-length cells will make for fast, efficient switching.

The general (but not universal) consensus regarding the architecture of high-performance telecommunications networks stops at Sonet and ATM. Left unanswered are such questions as: how will these networks support the interconnection of local-area networks? Will ATM be used in the local area, creating a seamless hierarchy? How will metropolitan-area networks be constructed? Will Sonet be appropriate all the way to the curb? To the desktop? What scheme will be used to provide pay-per-use service to the customer?

**WORLDWIDE RESEARCH.** Telecommunications carriers around the globe are experimenting with various schemes that may help answer these questions. Some are using a combination of Sonet and ATM to implement services like frame relay and switched multi-megabit data service (SMDS), both of which offer pay-per-use services as an alternative to leased 1.5- or 45-Mb/s circuits.

In Japan, Nippon Telegraph and Telephone Corp. (NTT), Tokyo, has supported narrow-band integrated-services digital network (N-ISDN) services since as early as 1988, making dial-up circuits at between 64 kb/s and 1.5 Mb/s commercially available. By 1996, NTT intends to support broadband ISDN (B-ISDN) services as high as 622 Mb/s commercially and to make available switching and transmission equipment with capacities of 10 Gb/s.

Advanced network research in Europe has included broadband (140 Mb/s) optical-fiber testbeds in Germany and FDDI-based metropolitan-area network services in Finland. At lower speeds, many European carriers are looking to frame relay to support switched 2-Mb/s services as well. N-ISDN is ubiquitous in France and is being deployed rapidly in other countries as well.

Two U.S. agencies, the National Science Foundation (NSF), Washington, D.C., and the Defense Advanced Research Projects Agency (Darpa), Arlington, Va., are behind an initiative to determine possible network architectures and applications for wide-area network services in the range of 1–3 Gb/s. Coordinated by the Corporation for National Research Initiatives (CNRI), Reston, Va., this effort asks for funding and/or collaboration not only from the Federal government but from industry and academia as well. A three-year pilot effort that began officially in early 1989, the project has five collaborative research testbeds—Aurora, Blanca, Casa, Nectar, and Vistanet—each with a unique mix of network and applications research.

All of the testbeds are investigating wide-

area network architectures, though geographically two testbeds (Nectar and Vistanet) are implemented over relatively short distances (within tens of kilometers).

Four of the five CNRI testbeds employ ATM technology, albeit each with different switch technologies and network architectures. The fifth (Casa) will use Sonet as a point-to-point transmission facility between Hippi crossbar switches. Three of the testbeds will use Hippi LAN and interface technology, and one of these three (Blanca) will also investigate Fibre Channel.

The other two testbeds are investigating the use of ATM host interfaces (Aurora), an IBM metropolitan-area network called Orbit (Aurora), and the Nectar host interface and LAN technologies of Pittsburgh-based Carnegie Mellon University (Nectar). Besides investigating ATM switching, Aurora will deploy a switch architecture from IBM called PlaNET, which is architecturally compatible with the IBM Orbit MAN.

**STORAGE TECHNOLOGIES.** Most mass storage systems today are based on a central server model, but advances in networks and storage technology argue for a change. Under the central server model, disk, tape, and other storage devices are shared among the machines on a local-area network by a central processor, which controls and grants access to the storage pool. Since that processor has to handle all data flow into and out of the storage devices, it must have a very high I/O capability—like that available only on large, expensive computers (mainframes and supercomputers).

Today storage devices can be attached directly to the network, and software is being developed to allow them to be shared by multiple computers. Better yet, networks can now provide capacity and bandwidth that is

## Network requirements are determined not only by processing speed but also by main memory size

in most cases higher than what is available on host- or network-attached storage media.

Among the newer reference models for mass storage systems was one proposed at the 1984 Sixth IEEE Symposium on Mass Storage Systems in Monterey, Calif. That model, which is still evolving within the IEEE Storage Systems Standards Working Group, provides a common framework for discussing mass storage approaches, components, and performance. A number of technologies and distributed mass storage systems are being developed in the context of this model.

As for software, a new low-level protocol that specifies a command set for networked peripherals now exists. The Intelligent Peripheral Interface 3 protocol, or IPI-3, as it is known, lets a host use a single network I/O channel such as Hippi to access multiple networked peripheral devices. Software above this level, based on the IEEE Mass Storage Reference Model, will permit multiple hosts to share the same networked peripherals by using a single control host as a traffic cop between hosts and peripherals.

Freed of the need for a central host to support both control and data flow, the mass storage system model for the '90s will consist of a set of hosts and devices, with one or more control hosts; a high-speed network for direct client-to-storage-device data flow; and high-speed, high-capacity storage devices directly on the network. Under this model, data flowing between the storage devices and the clients bypasses the mass storage control hosts, which therefore need no longer be large-scale mainframes.

The distributed nature of the IEEE Mass Storage Reference Model also makes separate hosts possible for distinct functions of the archive, such as name service (mapping file name to physical location), disk service, and tape service. In addition, it permits multiple disk servers and tape servers to work together.

That means that multiple disk servers for remote sites and workstation networks can be coupled with a central, high-performance system of networked disk and tape. The central, high-performance storage environment will be used by supercomputers and will also serve as secondary storage for the remote and workstation network servers.

The prospects of teraflops computing are enormous. To be fully realized, however, they will have to be supported by networks, storage devices, and archives of suitable capability. Otherwise, only a few applications—those that run on a single computer and do not generate too much data—will be able to use these magnificent machines.

**ABOUT THE AUTHOR.** Charles E. Catlett is associate director for computing and communications at the National Center for Supercomputing Applications (NCSA), University of Illinois at Urbana-Champaign. He is principal investigator for the center's work on the Blanca Gigabit Testbed.

**ACKNOWLEDGEMENTS.** Much of the work at NCSA described in this paper is supported by a grant from the Corporation for National Research Initiatives as part of the Blanca Gigabit Testbed. NCSA is one of four national high-performance computing centers supported by the National Science Foundation. It is also supported by the State of Illinois, the University of Illinois, and industry. Special thanks to Don Tolmie (Los Alamos National Laboratory), Randy Butler (NCSA), and Jim Hughes (Network Systems Corp.) for their comments and suggestions. ♦



# Modeling reality

*As supercomputers open new worlds to scientific enquiry, they are taking on Nobelist Wilson's Grand Challenges as well as mundane tasks*



Supercomputers rank right up there with the telescope, the microscope, and the atom smasher. They, too, are a technological advance that opens up new worlds to science. Mathematical models running on these superfast machines generate almost exact representations of reality,

which researchers use to study in unprecedented detail the phenomena of chemistry, physics, cosmology, biology, medicine, and other sciences, as well as some till now inaccessible areas of interest to humans.

National-security applications, in code breaking and nuclear-weapon design, for instance, are another forte of the big machines. Those uses are officially secret, but everybody knows they gave the supercomputer industry its start.

**CHALLENGES GALORE.** Then, almost 10 years ago, Nobel physicist Kenneth Wilson suggested making a collection of projects that presented "Grand Challenges" to researchers and their supercomputers. The list is now very long, but some idea of what it contains may be gleaned from the following sample: designing aircraft; simulating semiconductor materials; analyzing fuel combustion; the rational design of drugs; understanding catalysis; designing protein structures; researching the human anatomy; imaging (as of Venus); predicting the weather; studying air pollution; ocean modeling; evaluating ozone depletion; and petroleum exploration.

These far-reaching endeavors are heady stuff. After all, the researchers applying supercomputers to their problems are doing what no one before them could do. In conversation with them, they speak glowingly of hoped-for breakthroughs in voices that are full of enthusiasm and excitement. The studies going on cut across every facet of life: the air we breathe, the weather, the

Morris Grossman Contributing Editor

ozone and global-warming problems, airplane design, drug research, improved therapies, more efficient engines and cars, new chemicals, stronger structural materials, faster semiconductors, and many other important areas of study.

And mundane problems are being addressed, as well. For example, a golf equipment company used a supercomputer for structural analyses of a new titanium-head driver so as to provide a larger "sweet spot" for hitting the ball better. A U.S. plumbing manufacturer applied computational fluid dynamics to simulate the flushing of a new toilet design. The goal was to develop a quieter toilet that used less water for the European market. And computer modeling on a supercomputer was used to simulate, and make more efficient, the process of filling bottles with liquid soap.

In fact, mathematical modeling has led many scientists out of the experimental laboratory to the computer room, where they can more quickly and cheaply address difficult challenges via mathematical algorithms. For instance, instead of relying solely on slow and expensive trial-and-error methods of design and wind-tunnel testing, the physical experimentation of aircraft design is now being bolstered by high-performance computing techniques, massively promoted by the Numerical Aerodynamic Simulation (NAS) Program run by the National Aeronautics and Space Adminis-

Mathematical modeling  
lets scientists do  
bolder experiments on  
a computer  
than in the laboratory

tration (NASA), Washington, D.C.

Just thumbing through a volume of the *NAS Technical Summaries* gives an idea of the huge number of studies that fall under the rubric of computational fluid dynamics. These activities are mostly unknown among the general public. Fluid dynamics touches on specialties such as astrophysics, aerodynamics, hypersonics, turbulence, chemistry, propulsion, and structural mechanics. The heavy use made of color graphics in



these summaries is explained on the publication's inside cover: "Color graphics is a powerful and flexible visualization tool for the science of computational fluid dynamics. Without color graphics, the millions of numbers that form the supercomputer solutions to the studies contained in this report and on the cover would be virtually incomprehensible."

**TERAFLOPS NEEDED.** All this effort is going on with today's supercomputers, which are capable of billions of floating-point operations per second (gigaflops). But really to fulfill the aims of such study projects, the research community agrees that it will need computers with at least a trillion floating-point operations per second (teraflops). Such performance is now on the horizon, a supercomputer manufacturers vie with one another to deliver on these promises. Existing supercomputers fall far short of a teraflops on an actual problem. Still, although one or two manufacturers claim it as a theoretical peak, most say they could deliver a teraflops machine in two to three years. David J. Kuck, director of the Center for Supercomputing Research and Development at the University of Illinois, in Urbana, recently tested about 10 supercomputers and reported that none approached their claimed peak speed. The most powerful machines delivered to date are Intel Corp.'s Delta and NEC Corp.'s SX-3/44 with peak speeds of 20 and 22 gigaflops for 64-bit arithmetic.

To enhance U.S. leadership in developing more advanced supercomputers and to fund the large investments required, senior U.S. government, industrial, and academic scientists and engineers have cooperated in the federal High Performance Computing and Communications Initiatives (HPCCI). For the 1993 fiscal year, the program will invest \$803 million in four coordinated efforts. That is 23 percent more than the 1992 enacted level. Various Federal agencies will funnel the funds to laboratories, private industries, and academe, which will then run the programs.

**TUNNEL'S END?** "When faster, multiteraflops machines become available, they may even replace wind tunnels to a large extent," predicted Olaf O. Storaasli, research scientist at the NASA Langley Research Center, Hampton, Va. He said that such machines have become crucial to the design and evaluation of the structure, aerodynamics, and advanced propulsion of the new High-Speed Civil Transport (HSCT). In his opinion, the





Cray Research Inc.

[1] Flame-front propagation and velocity distribution in a four-valve gasoline engine [top] and flame-front propagation in a diesel engine [above] are visualized with Cray Research Inc.'s Multipurpose Graphic System working in a CRI/TurboKiva engine combustion-modeling environment.

mathematical modeling approach, besides being faster, more efficient, and less expensive, will lead to a supersonic plane that emits much less noise and exhaust emissions harmful to the upper atmosphere.

Fluid-dynamic codes running on supercomputers will, he said, speed the evaluation of alternative designs for aerodynamic efficiency and stability and the prediction of the strength and effect of associated sonic booms, the atmospheric heating of the vehicle, and the physical stresses of supersonic flight. The upshot will be the information needed to develop advanced materials and fabricate new structures, both to build military vehicles and to commercialize the HSCT, said Storaasli.

**PARALLEL RESCUE.** An approach that experts say will eventually lift supercomputers to teraflops speeds is the use of several processing units operating in parallel, instead of a single very high-speed processor operating serially. For years, building a fast computer meant increasing the processor's clock rate by the use of ever faster components.

The serial Cray-1 machines are a case in point. But at the start of the 1980s, the silicon-semiconductor emitter-follower technology used by the Cray-1s neared its inherent speed limit. Even packing components more closely and compensating ingeniously for the resulting heat increase had its limit, although upper bounds predicted in the '70s are now routinely exceeded.

Accordingly, many say that 1992 is the year when supercomputers are first defined by the number of parallel processors they can include. For example, the Cray-2 machines with four processors squeaks into the parallel category. But massively parallel computers, like the new CM-5 Connection Machine, made by Thinking Machines Corp., Cambridge, Mass., may prove a truer milestone in parallelism. It bridges the two approaches to parallelism developed in the '80s: the single-instruction, multiple-data (SIMD) processing of the older CM-2 machines and the multiple-instruction, multiple-data (MIMD) configuration of, say, the Cray Y-MP series.

Another possible milestone in massively parallel machines is the Touchstone from Intel Corp., Santa Clara, Calif. Storaasli and his associates use the Touchstone Delta Supercomputer, which has 512 parallel nodes, to develop parallel software for the structural analysis of a future supersonic plane. He explained that the Touchstone can handle the finite-element model of the HSCT, which "contains more than 14 000 structural nodes, involves 88 000 equations, and requires 3.3 Gbytes of memory to hold its stiffness matrix and geometric data." He claims that the new parallel algorithms can generate and assemble the HSCT matrix 66 times faster than can a single-head serial Cray Y-MP.

**INTERNAL COMBUSTION.** One of these Grand Challenges, analyzing fuel combustion in an internal-combustion engine, is by some measures a complex as the fluid dynamics of supersonic aircraft. It costs a lot in both time and money to optimize an engine design experimentally by varying all the design factors one at a time and testing performance in an actual machine. (This is one reason why leaner- and cooler-running machines are difficult to design.) In contrast, a computer model of the engine allows, at a fraction of the cost, quick parameter changes and simulated testing in hours rather than the weeks, months, or years it takes to build a physical model, fit it out with instruments, and try it out. In fact, today, the detailed simulation of what happens in an engine's combustion chamber has become almost routine.

Take the CRI/TurboKiva, a software environment that simulates engine combustion. Its promoters say it cuts back heavily on the period needed to test the combustion emissions and efficiency of new engines and allows users to quickly bring to market engines that are sounder environmentally and so have a distinct competitive edge. Working with the CRI/TurboKiva software, the Multipurpose Graphic System (MPGS) from Cray Research Inc., Eagan, Minn., meaningfully portrays the results of the modeling and helps in interpreting them [Fig. 1].

**CHIPPER CHIPS.** Of particular interest to electronics engineers is the work at Arizona State University in Tempe being directed by John D. Dow, a professor in the department of physics and chemistry. The university has installed a Convex C3340 supercomputer, a four-processor system with 256M bytes of memory, to simulate the growth of crystalline semiconductors (silicon and gallium arsenide). Dow's group is using the results to determine how to grow better and faster chips for the next generation of supercomputers. The Convex C3 series of supercomputers, priced from \$295 000 to \$8 million, may have up to eight processors, 4G bytes of physical memory, and a peak performance of 2 gigaflops.

Dow stressed how much advanced graphics, in his case, contributed to his analysis of the structures and properties of semicon-



ductor materials. "We compute moving pictures, but each frame takes millions of times more computing power than merely computing numbers," he told *IEEE Spectrum*. "With motion pictures at almost 1000 frames per minute, our work can exhaust any supercomputer in existence, and any likely to be built for at least a decade."

Even so, studying the videotapes helps his group design chips and optimize the chemistry of chip growth, and saves them the time and expense of building and running multimillion-dollar chemical reactors. To quote Dow further: "Every atom of a semiconductor crystal must be in its correct place. We are literally designing materials the way an architect designs a building—but one atom at a time." The arrangement of the atoms strongly affects the speed of the chip, and only a few atoms out of a few billion may disable the whole chip, he observed.


And just in July, Cray Research announced it was working on a "first of its kind" software tool for electromagnetics simulation and design that will allow the study of how electromagnetic waves interact with the surfaces of various objects. With this capability, the new ElectroMagnetic Design System (EMDS) can be used to model structural materials, such as those used for aircraft and missiles, that are being designed for reduced radar reflectivity.

According to Cray, aerospace engineers can use EMDS simulations to see how electromagnetic waves impact and scatter from the vehicles' surfaces, a capability never available before on supercomputers. The software is being tested at customer sites and Cray said it will be available on its Y-MP product line later this year.

Future releases of EMDS will provide for modeling the behavior of electromagnetic waves at the electron level. This capability will prove valuable in the design of leading-edge integrated circuits, according to Cray Research. The software is already part of a Cooperative Research and Development Agreement between Cray and the Department of Energy's Los Alamos Laboratory in New Mexico, which focuses on developing advanced software to simulate electromagnetic wave effects in ultrahigh-speed electronic components.

"The industry needs methods for directly solving the equations that govern the behavior of electrons in complex, nonuniform chip materials," said Los Alamos director Sig Hecker. "EMDS is an excellent basis for developing these important methods." Los Alamos is set to be the first test site to install the design software.

**DOCTORS' DILEMMA.** Far more complex than semiconductors are organic materials. The National Institutes of Health (NIH), Bethesda, Md., part of the Health and Human Services (HHS) agency, is the world's largest biomedical research institution, according to Robert Martino, chief of its computational sciences and engineering division.



Cray Research Inc.

[2] Images rendered by Tom Palmer of Cray Research Inc. and Frederick H. Hausheer, formerly with the University of Texas Health Science Center, now with his own firm, show how an interstrand DNA crosslinks in the body with phosphoramidate mustard, so that a bond is formed with the cancer cell's DNA that prevents its replication, thus killing the cell.



Jet Propulsion Laboratory

[3] Under the direction of Terry Cole of the Jet Propulsion Laboratory, a helicopter view, in three dimensions, of a portion of the Venus' surface was rendered on an Intel Touchstone 512-node Delta parallel supercomputer.

It operates an Intel parallel supercomputer to study the structure of viruses such as those causing herpes and acquired immunodeficiency syndrome, or AIDS. And the machine also models brain functions. Meanwhile, the agency is evaluating advanced supercomputer architecture and system software for use in biomedical research.

As for software aspects, the NIH concentrates on the algorithms needed for new database methods; structure determination from X-ray and magnetic-resonance imaging (MRI) data; studying and comparing amino-acid molecular-sequences in proteins and ribo- and deoxyribonucleic acids (RNA and DNA); and predicting biological structure and function from genetic code. The HHS program's analysis of normal and disease-associated genes complements the human genome project, which studies the genetic makeup of the human being.

In a very important NIH/HHS series of applications, supercomputers will rationalize drug design for such diseases as cancer and AIDS by improving the understanding

of molecular structure and function. Till now, trial and error has predominated.

For example, "development of new anticancer agents has been largely by trial and error efforts, and more recently by in-vitro preclinical screening with various malignant-cell lines and then human clinical experimental trials," said Frederick H. Hausheer, M.D. (He was formerly with the University of Texas Health Science Center, in the Department of Medical Oncology, at San Antonio, Texas, and now has his own firm, BioNumerik Pharmaceutical L.C., in the same city.) He noted severe limitations to that approach, because the projected screening of 10 000 to 15 000 compounds per year produces one potentially useful anticancer agent every two to four years and a wait of about 12 years from the first synthesis to a clinically useful substance. According to Hausheer, at present there are about 45 approved anticancer agents, and perhaps 10 of these are really helpful.

Hausheer hopes that supercomputers will substantially speed the process. For in-



stance, one area involves the use of *ab initio*, or very basic, fundamentals to study the chemical, conformational, electronic, and energetic properties of various clinically active drugs. The most commonly used agent for treating many forms of cancer (solid tumors, leukemias, and lymphomas) is cyclophosphamide. The active metabolite of that drug, phosphoramidate mustard (PM), is the anticancer agent [Fig.2].

The exact mechanism by which PM determines which cells are cancerous is only partially understood. In the body, PM gives up two chlorine atoms in exchange for two nitrogen atoms, forming a bond with the cancer cell's DNA and preventing its replication, thus ensuring its death.

Unfortunately, when PM is finally broken down by the body, it releases a toxic component. To learn how to overcome this effect requires a better grasp of how PM (and also other anticancer drugs) work. But that is, in part, beyond present computing capability.

Moreover, most biochemical reactions proceed in fractions of a microsecond, whereas current calculations simulate only fractions of nanoseconds and are thus thousands of times short of simulating interesting kinetic reactions. Clearly, handling all these factors by numerical methods calls for larger memories, faster computation (in the teraflops range), and advanced software—all still in the future. Hausheer believes that a massively parallel computational capability is required to implement such work.

Another biomedical example involves AIDS. The structure of the human immunodeficiency virus (HIV) includes protease, an enzyme that provides a necessary step for the maturation of the virus, enabling it to infect human cells. A supercomputer has been able to identify new agents that may block the action of the protease. Supercomputer studies have been able to identify several potential anti-HIV agents but, with PM, still more computing capability is needed to be able to chemically alter these new agents in a rational manner, so as to improve their therapeutic potential while reducing unwanted side effects.

On a more fundamental level, many biological processes are catalytically controlled by enzymes, which help select a desired reaction from competing reactions. Moreover, one-third of all industrial chemical processes use catalysts in chemical reactions, and catalysts are also fundamental in medical science, food processing, environmental restoration, and so on. Yet most are arrived at empirically and without much understanding of how they work. So now supercomputers, along with special algorithms, are being used to create new catalysts and further their understanding.

**IMAGING MARCHES ON.** An ever-present help in catalyst development and many other problems is the imaging algorithm. With its

aid, the internal structures and external shapes of molecules and macro objects such as the human anatomy may be displayed, rotated, viewed from any angle, and reversibly dissected. These transformations are very computer intensive, and the data sets obtained are large. For example, the Visible Human project will generate a 4-terabyte image library of human anatomy.

Other algorithms assist in simulating the collision of subatomic particles, studying the cosmological evolution of the universe, and modeling geological interactions—large-scale mantle convection currents, continental plate movements, and the resulting biogeological events.

In a geological simulation, engineers from British Petroleum Exploration Inc. (BPX) recently studied how to predict the amount that could ultimately be recovered from an underground oil reservoir, the rate of production, and the impact of operating decisions. Using a Cray Y-MP 8E system, the engineers were able to simulate a 1.5-million-grid block model involving sequences of hydrocarbon productive sands mixed with nonproductive shales.

The simulation of a whole reservoir (as opposed to a piecemeal approach) for a 10-year period ran in six and a half hours. BPX engineers said that 50 percent of engineering time can be saved on even more direct modeling methods and finer-scale numerical studies. Potential savings would be over \$7 million per year with better accuracy in predicting a well's oil production.

In another case, imaging the planet Venus's surface from data collected by the Magellan spacecraft since Sept. 15, 1990, needed a supercomputer for a realistic outcome. The Magellan used radar to scan over 90 percent of Venus for surface-structural

## Drug design by trial and error is slow, but supercomputers can rationalize the process

information, which it sent to earth in over  $3 \times 10^{12}$  bits (3 Tb) of data. This huge volume of data necessitated new and ever more labyrinthine methods of handling, processing, and storing it.

By way of comparison, generating just a single static image on a high-performance workstation's display can take over an hour. Even a short animated video version of a small stretch of Venus could take weeks to complete. However, an animated helicopter video rendering (viewable on a videocassette recorder), in three dimensions, of a portion of the Venus surface was generated

on an Intel Touchstone 512-node Delta parallel supercomputer in days [Fig. 3]. The effort was directed by Terry Cole of the Jet Propulsion Laboratory, Pasadena, Calif.

**SKY, SEA, LAND.** Yet another application area for supercomputers falls under the purview of the National Oceanic and Atmospheric Administration in the Department of Commerce, headquartered in Washington, D.C. The agency formulates and carries out operational and research programs in weather prediction and ocean sciences in collaboration with other HPCCI programs, such as those conducted by the National Center for Atmospheric Research, the Climate and Global Change Research Program, and the Coastal Oceans Program.

Modeling the weather has changed markedly. In 1982, to help visualize wind and precipitation conditions, P. S. Ray at Florida State University in Tallahassee and R. B. Wilhelmson at the University of Illinois in Urbana-Champaign strung together a rather crude physical model out of wire and layers of Plexiglas. Three years later, J. B. Klemp, headquartered at the National Center for Atmospheric Research (NCAR) in Boulder, Colo., put together a much more sophisticated model of the weather, but in a computer with color graphics to display static storm-cloud images.

Working with a Cray 2 supercomputer at the Colorado center, Wilhelmson two years later created a short video movie of the turbulent evolution of a storm for a single set of parameters. He told *Spectrum*, "My associates and I spent over a year of tedious work making the video, despite the Cray 2's ability to handle almost 2 gigaflops. With a teraflops capability, we could make such a video and then change the parameters and run it again in a few hours."

The ocean is the other turbulent fluid in the earth's climatic system. Its currents are much smaller and slower to respond than atmospheric movements, though they transport almost as much heat. Ocean currents on average are only about 50 km wide, and ocean eddies are a tenth to a hundredth the size of those in the atmosphere.

The first global ocean model can barely resolve eddies and strong currents, and can simulate conditions for only about a 10-year span. It has only recently been completed by the National Oceanic and Atmospheric Administration. The work required 3000 processor-hours on a gigaflops Cray-2. The model produced more than 250 gigabytes of processed output. A two-hour animation video was made to interpret the accumulated data, requiring the transfer of gigabytes of model output across high-speed networks. But to simulate a longer period, such as the several centuries of ocean activity needed really to predict oceanic events, would require a teraflops supercomputer.

Other influences vital to this planet's well-being—but not well-understood—occur be-



yond the atmosphere, in the stratosphere. The recent ozone losses there are being intensively studied. But the chemical and dynamic mechanisms that control the depletion are extremely complex. Hundreds of chemical processes are involved, as well as the details of mixing and movement of the air containing the chemical compounds. According to Robert Ward of the Oak Ridge National Laboratory (ORNL) in Tennessee, today's ozone-depletion models need as many as 10 supercomputer hours to simulate only one day of the activity, but years of activity must be simulated to fully understand the trends.

**MODELING TURBULENCE.** Solving the problems of smooth, laminar flow in fluids like water or air is made relatively easy by well-established methods. But it is not practical to handle turbulent flow around the wings of new airplane designs, say, with pencil and paper applied to the Navier/Stokes differential equations. Indeed, the Navier/Stokes equations still apply, but their full solution requires more powerful computers than are presently to be had. A massively parallel supercomputer architecture appears to be exactly what is needed.

A rather interesting programming approach to turbulent-flow study, eminently suited for parallel architecture, was developed by a team of scientists at the Los Alamos National Laboratory, New Mexico—Uriel Frisch, Brosl Hasslacher, and Yves Pomeau—plus Jim Salem, Bruce Nenich, and others from Thinking Machines. Instead of using standard numerical methods with accurate real-number arithmetic to simulate a process, the team's approach follows the movement of a large number of individual particles, emulating the movement according to the kinetic theory of gases and fluids.

In a process similar to a Monte Carlo procedure, each particle is allotted some initial direction (but the same average speed, to simplify the initial conditions), in line with how the molecules in a fluid might behave. The particles (molecules of the fluid) then start to move and interact in a large grid of cellular lattices in conformity with Newton's laws of the conservation of momentum. Accordingly, the vector sum of particles moving inside each lattice, before and after each time step, is the same. Up to six particles can collide at once from six different directions, hence the name of hexagonal lattice.

The grid was created by configuring the Thinking Machines' CM-5 as a plane with links between each processor and its four immediate neighbors (although each processor may in this particular machine have links to 12 neighbors, allowing the CM-5 to act as a hypercube in 12 dimensions). This planar arrangement keeps the process moving at the maximum time-stepping capacity of the machine. Each time step consists of two parts: a communication and a computational phase. After each time step, the computer checks for the condition of any collid-

ing particles. The total number of possible types of collision combinations is 64, but many are redundant, so that by proper manipulation the number is reduced to a more manageable 14, the number of rules the working model in fact contains.

But complete solutions to problems involving turbulence with either Navier/Stokes or the lattice gas dynamics methods are not yet practical with serial supercomputers because of the huge number of separate and independent calculations needed for the analysis. Nevertheless, for now, designers are busy reducing wind-tunnel testing with partial solutions obtained mostly on serial computers.

**HOW FAST IS FAST?** Speed trials on several competing machines were conducted about two years ago by the Center for Supercomputing Research and Development at the University of Illinois in Urbana-Champaign. The benchmark software used was Performance Evaluation for Cost-Effective Transformation (called Perfect). An eight-processor Cray Y-MP/832 costing about US \$19 million took top honors, with 120 megaflops for a hand-tuned, highly optimized program, and just 22 megaflops untuned. Slower, but at less than 1 percent of the cost (\$123 000), a Stardent-3000 superworkstation without tuning delivered nearly 20 percent of the Cray's speed, or 4.2 megaflops. Thus, the biggest machines are often not the most cost-effective—or even available when needed [see also pp. 39–41].

For example, researchers get free computer time at the four supercomputing centers run by the National Science Foundation at the universities of Pittsburgh, Illinois, San Diego, and Cornell. Each has a Cray machine, and together the four have somewhat more capacity than 30 Cray X-MPs, but as many as 10 000 researchers vie

An oil well's production  
can be modeled on  
a supercomputer  
so accurately as to save  
millions in overhead

for access to them. The machines represent about 240 000 processor-hours per year, or, on average, little more than 24 hours per researcher per year. A high-performance PC delivers that in a month, which may, however, be enough for the occasional user. As a result, only 5–10 percent of the researchers use the four centers. Most buy smaller machines.

Moreover, seemingly comparable supercomputers ran the same program at different rates, some several times faster than others, and some approached its touted peak.

Such a wide range of performance bothers David J. Kuck, director of the Urbana center. "No one outside the research community has the patience to cope with such unpredictable behavior," he said. "Instead, commercial users want the predictability of machines such as the VAX 11-780 and DEC 6000 series. The supercomputer manufacturers have only one answer: 'True, but ours are difficult machines.'"

**WHO NEEDS ONE?** On some small projects, less costly machine types will do instead of the top supercomputers. Superworkstations combine varying degrees of supercomputing capability to provide 10–30 percent of the Cray Y-MP at often less than 2 percent of the cost, as their prices range from \$50 000 to \$200 000. Next in cost-effectiveness come minisupercomputers that provide up to 20 percent the capability at 16 percent of the Cray Y-MP's cost. Mainframes, however, with up to 25 percent of peak Cray Y-MP capability and as little as 50 percent of the Cray's cost, are not highly cost effective, since many of them sell for more than some of the low-end Crays.

Clearly, the best of these substitutes for the top supercomputers are superworkstations. Not only do they provide a high performance-price ratio, they also interface with users more easily. They can be linked in a parallel network, which may outperform even a Cray Y-MP on some problems. Moreover, being based on fast-evolving technologies, they can be expected to improve in performance by 50 percent per year, according to some experts. Also, expensive machines like supercomputers or mainframes can involve long delays due to exhaustive bureaucratic shuffling over financing their acquisition.


For now, some feel that the serial supercomputer is a protected species because of its use in the Department of Defense and because it has become a symbol of the United States' superiority in computer technology. Others say that there is no "conspiracy" in their favor, just that new parallel technology takes time to develop. Parallel machines have only a small installed base so far, but the expectation is that by 1995, massively parallel computers will have the trillion flops needed by so many programs. But who can predict? By the year 2000, PCs may have the capability of today's supercomputers.

**ABOUT THE AUTHOR.** Morris Grossman, an electronics engineer and registered professional engineer, has over 30 years of experience in the field. He has worked at Warner & Swasey Inc., RCA Corp., and Control Data Corp. Now a technical journalist, Grossman was also an editor for 17 years at *Electronic Design* and *Electronic Products*. A previous article of his for *IEEE Spectrum* was on information theory and was titled "On the nature of information—an application of entropy" [December 1965, pp. 70–72 and 77–80].



# Perspectives on visualization

*After bouncing between researchers and entertainers, this embryonic field is coming into its own—and the payoffs could be immense*



Akira Kurosawa's movie *Rashomon* explores an event—a man's murder and the rape of his wife by a bandit—through the eyes of the bandit, the wife, and the husband. The three perspectives differ wildly, and while none is the "truth," all contribute to the viewer's understanding

of what happened.

Clearly, the understanding of an event is mediated by the point of view, and it is possible to discriminate between the perspective and the event as we observe what transpires from two or more points of view.

In the field of computational research, events are abstractions, which may or may not be based on physical phenomena. The scientist can gain access to these events through theory, mathematics, computer code, numerical data and graphics, and other techniques. The more techniques applied, the more perspectives are brought to bear on the problem.

Visualization is commonly associated with three-dimensional graphics, but it applies more generally to the use of visual forms to represent information to the viewer. In most cases, the visual representation has been found to be more intuitive or accessible than numerical or mathematical portrayals of the same data. But each of these representational forms—mathematical, numerical, and visual—has its pros and cons in conveying information.

Comprehension of information emerges from the interplay and the juxtaposition of these various representational forms. The juxtaposition of visual, textual, and numerical forms, for example, gives the viewer graduated access to the data, guiding him or her from the qualitative to the quantitative.

Matthew Arrott

National Center for Supercomputing Applications  
Sara Latta Science Writer

Many of the most interesting aspects of visualization stem from how the various forms are combined and juxtaposed. It is here that visualization gives way to the larger issues of information design, the field devoted to effectively and succinctly explaining and representing concepts that are often complex and vital.

**TRAGIC OMISSION.** For example, the night before the space shuttle Challenger accident in 1986, engineers from Morton Thiokol (the makers of the shuttle's solid rocket motors, which boost the shuttle into space), the National Aeronautics and Space Administration's Marshall Space Flight Center, and the Kennedy Space Center discussed the  $-0.5^{\circ}\text{C}$  temperature forecast for the next morning's launch. Because the lowest previous launch temperature had been  $11.6^{\circ}\text{C}$ , they were concerned that the low temperature might affect the performance of the engine's O-rings sealing the field joints of the solid rocket motors.

They used a graph [p. 65] to analyze the historical data correlating space shuttle incidents (thermal distress of the O-rings) with temperature. Because of the U shape of the graph of thermal distress incidents versus temperature, they concluded there was no evidence for a temperature effect on O-rings. They were, of course, fatally wrong, and their error lay in part in not plotting the flights with no thermal distress incidents.

Graphs, images,  
and animation make  
sense out of masses  
of otherwise  
incomprehensible data

The commission appointed by President Reagan to investigate the cause of the accident noted that a "mistake in the analysis of the thermal distress data... was that the flights with zero incidents were left off the plot because it was felt that these flights did not contribute any information about the temperature effect," according to a report in the *Journal of the American Statistical Association*. "The commission concluded that 'a careful analysis of the flight history

of O-ring performance would have revealed the correlation of O-ring damage in low temperature.'"

**THEORY TO REPORTING.** Computational research is done in phases: theory and mathematical development come first, then algorithm development, computing analysis, data and computational management, and finally reporting of results. Visualization started as a tool for analysis and reporting, but users quickly discovered that it was useful for some of the other phases—computing and data management, and code debugging, for example. Lately it seems to have become so all-encompassing that we have at times lost sight of the fact that it is only a mechanism for conveying information compactly and effectively.

As an example of visualization's role in a computational problem, consider the formation and evolution of a severe thundercloud. The atmospheric scientist collects such information ■ wind speed, air pressure, temperature, water vapor, and cloud and rain water content [Fig. 1, next page]. The collected information is used to create theories of behavior. From these theories the scientist can derive a mathematical model [Fig. 2]. The mathematical model is turned into computer code, which in turn can produce a deluge of data [Fig. 3]. Traditionally, visualization made this mass of data comprehensible by producing graphical images—in this instance, an image of a three-dimensional form [Fig. 4].

The process of finding a fitting visual carrier for scientific information is the same as for other types of data, whether they are two-dimensional portrayals of ■ integrated circuit layer or three-dimensional representations of a management structure.

Perhaps the most familiar of these idioms is the three-dimensional contour. In the case of the storm cloud model, the cloud water and rain water are what reflect light and give the cloud its visible form [Fig. 5]. Color may be

used either to distinguish between the cloud water and rain water, or to illustrate different directions of movement.

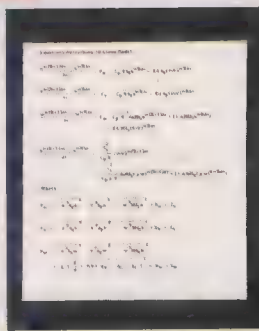
**2-D WOES.** Of course, representing three-dimensional forms on the two-dimensional plane of a monitor screen inevitably obscures part of the information, so a number of techniques have been developed to cope with this problem. If one surface is nested inside another, the scientist can reveal the inner surface by making the outer one trans-




## Evolution of a thundercloud




1




2



3




4



5



6



7

parent. To see the whole object, the computer may keep the object fixed and move the viewer around it, or keep the viewer fixed and rotate the object.

Also, the scientist may take a two-dimensional slice out of a three-dimensional geometry, and by putting the two forms side by side, provide context for the two-dimensional plane and detail for the three-dimensional geometry [Fig. 6]. A one-dimensional color bar supplies the necessary quantitative guide to the use of color on the plane.

Trajectories, or historical traces, which show how particles move, are another method of assisting the viewer and giving the illusion of three dimensions. Too many of them make the image look like a bowl of spaghetti, however, so sometimes they are used in conjunction with particles to show location and movement. In the storm cloud model, which is animated, particles show the location of the updraft, and the motion conveys the dynamics of the particle flow [Fig. 7].

Even with these graphics conventions, information is still lost in translating a three-dimensional scene into a two-dimensional image plane. To compensate, there are other tools and techniques, which generally involve the rearrangement of color on a two-dimensional plane. Some of these techniques reinforce the original information, giving the *appearance* of three-dimensions.

A silhouette of a quasi-oval shape, for example, tells the viewer no more about the object than a line drawing [Fig. 8]. The object might be an egg or a stone, but then again it might be a faraway moon or planet. Drawing a grid on the shape [Fig. 9] conveys the shape of the object, but no information

about its interior or the scaling of the grid.

Because people are used to sizing up and comprehending three-dimensional shapes from the way light washes over them, adding light to the scene gives additional information—about which areas of the form are exposed to it, for example [Fig. 10]. A highlight reveals more about the object in the context of the light and the viewer [Fig. 11]. All these techniques add to the viewer's understanding of the geometry.

People also perceive objects in relation to their surroundings. Thus a graphics programmer can establish a three-dimensional relationship by setting an object on a stage, or casting shadows on the stage [Fig. 12]. The image might be a composite of several images viewed from different angles, establishing the relationships between those perspectives with some key points within the scene, so that the mind can integrate and assimilate the information across the images [Fig. 13]. Motion, of course, introduces the fourth dimension, time: the three-dimensional object can be rotated to reveal parts of the form not seen earlier. Also, the viewer's vantage point can be shifted in relation to the whole scene in order to pick up the movement of one object with respect to another.

Certain visual idioms may be aided by the reinforcement of other sensory representations. One is sonification, the representation of data by sound. Cardiologists analyzing a model of the blood flow throughout the body, for instance, learn from an auditory as well as visual representation of the blood flow. Physicians have long used stethoscopes, so sonification is an ideal carrier for this type of scientific data.

A number of media forms can display visual images. A still photograph lets the viewer study or peruse the information within the image at will. Because they do not change, however, they limit the viewer to a single perspective and moment in time. An image sequence—an animation or motion picture—gives the viewer the dimension of time, but it deprives him or her of the ability to explore a single image at will, and often fixes the rate at which the sequence unfolds.

**NEEDED: INTERACTIVITY.** What the viewer requires above all is interactivity, the kind offered by a graphics workstation. Interactivity gives the user options in conveying information: visual, auditory, numerical, or mathematical, for example. Processes may be represented as graphs or as lines of code. The viewer can juxtapose parameters to discover and map their relationships.

Simultaneously displaying representations of temperature and pressure, for example, provides a context for the effect those parameters have on an event. The viewer can also place a representation of a computational process next to one of the data it generates, to better understand the computation and the phenomena it models.

This brings up an inherent tension in the practice of visualization: it is expected to perform as a tool for analysis and exploration of data, as well as a tool for communication of that data to an audience. The viewer as scientist analyzes and explores the data, moving from a position of unfamiliarity to one of understanding; as such, the visualization tools are designed to be used for exploration.

But once the scientist grasps the problem by using visual tools, how can research



results be conveyed to peers and the scientific community at large? Visualization as a form of communication is the logical answer; unfortunately, most of the visualization software built to analyze data lacks the tools to communicate that same data.

There have been tools for analysis and others for communication, but until recently, there have been no packages designed to perform both functions. These limitations are disappearing at a reasonable rate. At the National Center for Supercomputing Applications (NCSA), for example, we are experimenting with integrating existing software technologies into an environment for studying atmospheric and fluid flows. This environment has both analytic and communicative capabilities.

**HISTORICAL SCHISM.** The dichotomy in the way visualization is used is traceable to the way the technology evolved. Historically, computer graphics have been used in two separate and quite different industries: for science and engineering, and for entertainment. The technical community is interested only in the visual image's ability to convey information about the data driving the image, and in most cases the images were very literal translations. So, for example, the electrical engineer designing a chip would produce an image of the layout of wires on the chip. In entertainment, on the other hand, the images themselves are the product, and the goal is simply to engage the audience.

In the early 1980s, the entertainment industry recognized computer graphics' potential for creating special effects in film and commercials and appropriated the fledgling field, then under development in academia. Walt Disney Co., Burbank, Calif., made extensive use of graphics in the movie *Tron*;

Digital Productions used them in *The Last Star Fighter*. Robert Abel and Associates used computer graphics in their TRW commercials.

In the mid- to late-1980s, the scientists and engineers reappropriated the computer graphics techniques developed by the entertainment crowd and reunited them with various ongoing university efforts. The technical community's goal was finding new ways to communicate, as well as analyze and explore, scientific data.

**MACHINES OVERWHELMED.** Today, the boundaries of visualization coincide with the limits of technology. To date, our desire for functionality has outrun the capabilities of the machines, as illustrated by an apocryphal story from the entertainment industry.

When the television industry began using computer graphics, a graphics expert spoke to a group of broadcasters. Someone in the audience asked how long it took to render a frame using computer technology. The expert replied: 20 minutes. The questioner persisted: what if I use a more powerful computer, or get faster software? The answer remained: 20 minutes. A faster computer or better software would simply allow the animators to raise the level of sophistication, within the constraint of keeping the time down to the usual 20 minutes. In other words, far more computer power would be absorbed in improving image quality and sophistication before the time to render frames started to decline.

Today, the level of sophistication that took 10 or 20 minutes in the mid-'80s can be achieved in a fraction of a second. We now have the computer capability that comes close to what we want, and we can finally begin thinking about what we can achieve

using advanced visualization techniques, ■■ opposed to concentrating on the mechanics of the process.

**BORROWED TECHNIQUES.** The point of entertainment visualization is quite different from that of technical visualization, but that has not kept engineers and scientists from capitalizing on the visualization work originally done for entertainment or graphics design. Perhaps the best example of this cross-fertilization is the construction of nonliteral representations of technical or scientific data.

When confronted with complex concepts or phenomena, people often try to understand them by creating pictures in their mind's eye, and engineers and scientists are no exception. Physicists speak of gluons, which "glue" quarks together in subatomic particles. "Voltage-gated channels" is the name cellular biologists have given to cell membrane proteins that allow small molecules to pass in and out of cells in response to a change in membrane potential. Scientists involved with visualization also use these visual idioms to get their points across.

Rustling leaves on a tree or the darting wisps of smoke from a chimney are the way we "see" the wind, for instance. Thus a stream of moving particles is a logical choice for showing wind speed and direction in an animation [Fig. 7].

Ideally, the process of assessing scientific information should be transparent for users who merely want to move information from the numeric to the visual domain. At some point, however, the user may need to modify the process, and here faces both technical and intellectual challenges. Moving the data from one machine to another,


## How light dispels doubt




8




11




9



12



10



13

National Center for Supercomputing Applications



or between applications, presents a technical problem. But more important is the intellectual problem of choosing the idiom appropriate for the information.

In adopting visualization as a tool for analysis and communication, scientists adapted it to their own needs. And just as people have adapted the principles of desktop publishing to many different fields, with different standards for users in different fields, so scientists have begun to develop standards for visualization appropriate to their various fields of study.

The developers of desktop publishing tools—who are after all not writers, editors, or designers—have significantly changed the software to meet the needs of the different markets they serve. So, too, the developers of visualization software are now modifying their tools as they gain a sense of the complexity of an application and the needs of its users.

One of the most significant recent achievements here is tiered user access—the accommodation of different levels of user expertise and intention. Using this feature, an atmospheric scientist, inexperienced in the use of visualization systems, might begin exploring a storm cloud model by simply connecting the computer simulation to the visualization and the analysis tools, and ob-

serving the results using the default settings of the program.

At some point, though, that scientist will want (or need) to change some of the parameters: alter some of the input data, juxtapose representations not provided for in the default settings, or create a new representation altogether. Tiered user access allows the scientist to extend the exploration beyond the developer's *a priori*

## Interactive visualization offers the possibility of profoundly changing the way we assimilate and process new information

idea of the likely exploratory pathways. Access to exploration is based on the needs and desires of the researcher, not on confronting the scientist with all the system's capabilities at all times.

Visualization is by nature a process: it begins with the transformation of data into visual idioms, continues by using the visual idi-

oms as maps to other representational forms, and finishes by displaying the image or image sequence on a computer screen or perhaps through a virtual reality environment. As users move to more and more sophisticated tiers of access, they can visualize the computational process as well as the relationships among the data.

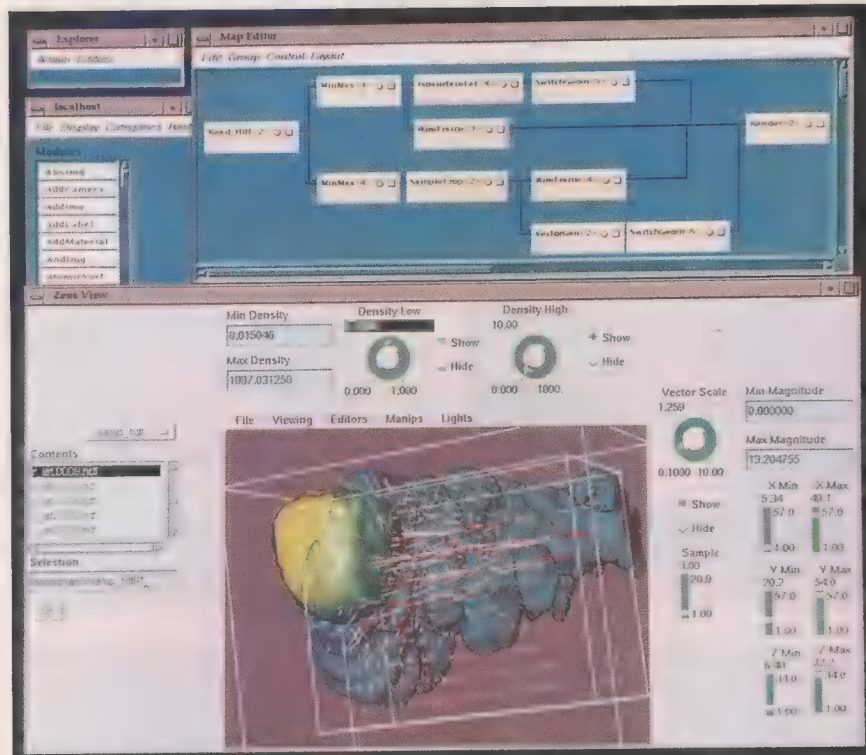
These tiers of access allow the user to navigate through and explore the process itself. Figure 14 is a visual representation of the process used to create a visual representation of an extragalactic radio jet colliding with an intergalactic cloud much denser than the jet itself. That is one part of the visual image. The other is the process organization, which is displayed on the same screen. Each box groups pieces of the application—a tool that renders the geometry of the image, a different one that reads files, still another that changes data, and so on.

If, after exploring the representation, the viewer realizes that a new parameter is needed, it can easily be added by modifying one of the tools. By placing the representation of the process next to the data, the user can better understand how changes to the process modify the data.

**FUTURE TRENDS.** As with many developing technologies, the pioneers in scientific visualization worked first to master the technical aspects—in this case, the mechanics of packaging data in a visual form. More recently, they have stepped back from the process to think about the implications of their work and what they are trying to achieve. How can visual, numerical, mathematical, and dynamic elements be best integrated to communicate or aid in analysis? How many representations can be juxtaposed without causing confusion? Could visualization free us from our age-old ways of understanding information? The answer to these kinds of questions will come from the field of information design.

As practitioners become adept and comfortable with visualization, they will begin approaching their problems from entirely new perspectives, and using pathways of logic that may not have been possible without the aid of the computer. After all, it is access to multiple perspectives that allows us to truly understand something, whether it is a problem in astrophysics or an orange. The orange is visibly round and orange-colored. It has a citrus odor. Peeling it reveals another layer. The juice inside feels wet and sticky; it tastes at once sweet and slightly sour. All of these perspectives together assure us we are dealing with an orange.

For much the same reason, interactive visualization could profoundly change how we assimilate and process new information. Ours is for the most part a linear information age. Our culture is strongly influenced by the written word, which is an inherently linear form of communication. The reader



Visualization is often used to illuminate phenomena too small, large, ephemeral, long-lasting, or far away to be viewed directly. In the screen above, various windows display all of the information used to create an image of the collision, in intergalactic space, of a radio jet and a much denser gaseous cloud. An image of the collision itself is visible at the bottom and center of the screen. Also displayed, in windows around the image, are all of the relevant data, code, and process organization that went into the visualization. Each box groups pieces of the application—a tool that renders the geometry of the image, a different one that reads files, still another that changes data, and so on.



follows the logic of the writer, word by word, paragraph by paragraph. Scientific papers state the problem, outline the methodology used to solve that problem, give the results of the experiments, and discuss the author's interpretation of them. The reader is forced to understand the paper through the author's logic, and indeed is seldom privy to other information: the paths of logic that seemed to come to a dead end, the failed experiments.

Such a plurality of perspectives could be accessed if the scientific information were "published" on compact-disc ROM through the use of interactive visual software. Such an alternative raises the possibility of an audience threading its way through a scientific experiment on its own, exploring different avenues. The reader could even become a collaborator of sorts, if he or she chose to become involved at that level.



This kind of interactive scientific communication may not be so very distant. Indeed, writers, the purveyors of that quintessentially linear form of communication, have already begun to write and publish fiction on hypertext software, a format that allows the reader to navigate the narrative at will. The story is different for every different reader, because the narrative logic is dictated not by the author but by the paths the reader chooses to take through the story.

**COMPLEX PERSPECTIVES.** An event cannot be understood from a single point of view. Scientific phenomena—and most other events—are too complex for all their elements to be grasped at once. We explore the event from many perspectives, viewing the elements on their own and in relationships. It is from this exploration that we create sophisticated perspectives in our minds.

Similarly, interactive analysis will allow the scientist to explore the data in ways not previously possible. Tiered access to the process of visualization will let scientists explore the data in ways that are quite new. Scientists will be able to change the calculations performed by the computer simulation, interacting with the simulation itself as well as its results. They will be able to change their focus within the volumes of information—placing the emphasis on one or another representation—as well as the process of representing the area of focus.

Scientists will also be free to choose the means of presentation—is the image best displayed on a computer screen in a linear set of sequences or in a hypermedia form, augmented by sound, or in a virtual reality environment? Interactive visualization better matches the nonlinear characteristics of scientific experimentation. It provides the means to present scientific experiment and discovery in a nonlinear form.

From its origins in the scientific, engineering, entertainment, and design communities, visualization has blossomed into an active, challenging endeavor in its own right, a component of the field of information design. Nonetheless, we still only know the process



Source: Journal of the American Statistical Association

*The night before the ill-fated launch of the space shuttle Challenger in 1986, engineers studied a graph [top] showing the number of times the shuttle rocket motors' O-rings had experienced thermal distress, plotted as a function of the temperature of the motor's field joints at launch time. The engineers concluded from the roughly U shape of the graph that there was no compelling evidence that temperature was a factor in the distress. However, the 17 flights in which there was no evidence of distress—all of which took place above 18.3 °C—were not plotted [above, in blue], and would have undermined the engineers' conclusion.*

of visualization from a few routes of investigation. To date, most of these have been focused on technology.

Efforts have been made to explore the use and implications of visualization, but we need a great deal more experience with the issues of representing information from different perspectives before we can gain a sense of what it really involves and what it has to offer. The glimpses of the possibilities are tantalizing, promising a day when engineers will fine-tune enormously complex designs with the help of powerful interactive graphics, and researchers will roam through intricate experiments at will, gaining critical insights that would be otherwise denied them. It does not require much imagination to see that the payoffs could be immense.

**ACKNOWLEDGMENTS.** The authors wish to acknowledge the contributions of the Visualization Group at the National Center for Supercomputing Applications at the University

of Illinois in Urbana. The Challenger example came from Edward Tufte's forthcoming book *Visual Explanations*.

**ABOUT THE AUTHORS.** Matthew Arrott has worked since 1987 on the development and understanding of scientific visualization at the National Center for Supercomputing Applications at the University of Illinois in Urbana. He directed the production of visualization videos including "The Study of Numerically Modelled Severe Storm" and "Smog: Visualizing the Components." Arrott received his B.A. from Harvard University in 1983.

After earning an M.S. in immunology from the University of Chicago, Sara Latta traded pipettes for pen (now personal computer). She writes on medicine, science, and technology for various publications, including the University of Chicago's *Medicine on the Midway*, the American Oil Chemists' Society's *Inform*, and the National Center for Supercomputing Applications' *access*. ♦



# Superchips for supercomputing

*The technology for fast gallium arsenide chips is making strides, but familiar silicon still has plenty of potential*



Semiconductor technology for central processing units of supercomputers is in a state of flux. Though long the leader, silicon emitter-coupled logic (ECL) now faces competition from silicon bipolar complementary metal-oxide-semiconductor (biCMOS) ICs and gallium arsenide ICs. In fact, a variant

of ECL, called emitter collector dotted logic (ECDL) may also vie with its progenitor [“A dotted kind of logic,” p. 68]. Even within the GaAs category, novel logic structures may soon jostle for position. To further confuse the picture, ordinary CMOS chips for workstations may be the best choice for massively parallel-processing supercomputers of the future.

In supercomputers' central processing units (CPUs), biCMOS is making a strong bid for fast RAMs, although it is also the primary high-speed logic in midrange supercomputers like Convex's C3400 family. In the present generation of fast RAMs, biCMOS usually has longer input and output delays than ECL (about 2.5 ns versus 1.2 ns for ECL). For one thing, the bipolar transistors produced by biCMOS processing are not as well optimized as ECL transistors. For another, signals to or from external ECL ICs have to be translated at biCMOS I/O boundaries.

I/O delays account for much of the total delay of small RAM chips, say 1K bytes, to the advantage of semiconductor processes with faster I/O. But as the amount of memory on a chip increases, the speed of the core memory itself plays an ever larger role in overall RAM performance. BiCMOS utilizes CMOS to implement a memory cell that is smaller than the ECL one, but, unlike pure CMOS, has bipolar devices to drive row and column lines and for sense circuitry.

So for large RAMs, to offset slower-than-

Harold Dozier Convex Computer Corp.

ECL I/O and performance reduced by bipolar-driven internal circuits, biCMOS offers smaller core memory, shorter internal signal paths (particularly on row and column lines), and overall faster core access. BiCMOS can be especially effective in self-timed RAMs, where internal core memory cycle time is the worst constraint.

For CPU logic, biCMOS uses a CMOS gate structure with bipolar-buffered outputs. The logic gates produced are smaller but slower than the ECL variety, but fast enough for midrange supercomputers. Representative cycle times are 20 ns for biCMOS versus 5 ns or less for ECL.

**CHIP TO CHIP.** Depending on the system architecture, part of the disparity in gate speed between biCMOS and ECL can be mitigated by the larger number of (smaller) gates per biCMOS chip. Extra functionality can be put on a biCMOS chip, to reduce the number of chip-to-chip delays.

As for GaAs versus ECL, until recently ECL held the edge in terms of available logic gates per chip (gate count). But GaAs is no longer immature. There is now near-parity in terms of performance and gate count between ECL devices and direct-coupled field-effect-transistor logic (DCFL) GaAs devices in volume production—and GaAs consumes less power.

GaAs also has more logic gates per unit area. ECL gate structures are not as simple nor as small as equivalent DCFL gate

power produces no advantage at the system level, the point is moot. But in the Convex C3800 supercomputer, air cooling, rather than liquid-refrigerant cooling, was a system requirement. Here, the lower dissipation of GaAs was the overriding reason for its being favored over ECL.

Incidentally, DCFL is not the only logic form that can be implemented in GaAs. The high-density GaAs chips now on the market are mostly DCFL, which has transistors much like silicon n-channel MOS transistors, but is usually built of only NOR gates rather than NOR, NAND, or complex gates.

This simple gate structure is what saves DCFL most power. Per function, DCFL requires less die area than does ECL. This means that signal paths within gates and especially between them are shorter than in ECL and are therefore more lightly loaded. Thus, the same switching delay can be achieved with lower current.

Current-mode logic (CML), which closely resembles ECL, can also be built in GaAs. With GaAs CML, faster-switching FETs replace the bipolar transistors in ECL, but CML mops up even more power than ECL. Moreover, GaAs CML structures are much more complicated than DCFL gates, so a GaAs-based CML chip would have a far lower gate count than either DCFL or ECL. CML is used in communications circuits, but seems not to be the primary logic technology in any commercial supercomputer. Still, that could change as GaAs technology advances.

The neglect of CML so far in supercomputers points up the fact that the machines' designers do not jump to use a new high-speed technology the minute someone announces that it has been fabricated in an R&D lab. It costs tens or even hundreds of millions of dollars to develop a new supercomputer, and companies cannot afford to have a project of that magnitude fail or slip because one of the parts used in the design turns out to be unmanufacturable.


Undoubtedly the supercomputer industry is driven toward high-performance technologies by its quest for ever-higher system performance. But just aligning a new supercomputer design with a new generation of higher-density, smaller-geometry ICs is risk enough; considering technologies that have no production history behind them is too dangerous. After all, it took the better part of a decade for GaAs ICs to advance from

The microprocessors used in workstations may be the best choice for massively parallel machines

structures, but ECL chips can be fabricated in larger sizes, at equal cost, thanks to the more mature silicon-based technology and the availability of larger silicon wafer sizes with lower defect density. The upshot is that ECL and DCFL GaAs are roughly equal in performance and cost. For both, a state-of-the-art clock speed is 200 MHz and gate count is 50 000 gates per chip.

Gallium arsenide's lower power dissipation can count for a lot. If dissipating less





Cray Research Inc.

*Module for a supercomputer consists of an array of ICs directly mounted on a printed-circuit board by gold-to-gold bonding. Each fully packaged IC contains 10 000 logic gates. Vertical pins are for interconnections to other circuit boards.*

a concept to a commercial product for supercomputers.

In short, supercomputers will probably continue to use the same three primary IC technologies—biCMOS, ECL, and GaAs—for several years to come. Advances are likely to come through improvements in these technologies rather than from commercialization of a new technology.

**MAYBE SUBOPTIMAL.** A case in point is biCMOS technology, which is ripe for improvement. Many suppliers created their biCMOS processes by adding additional—but not necessarily optimal—bipolar capability to an existing CMOS process. Others have added not necessarily optimal CMOS capability to an existing bipolar process. Some chip suppliers have even taken both approaches.

One view of the future holds that there will not be separate ECL, biCMOS, and CMOS processes, especially for logic and high-speed RAM, though separate, specialized processes will continue to be used for dynamic RAM. Semiconductor vendors will create a single process that can simultaneously produce very good, high-density CMOS and high-performance bipolar transistors and resistors.

Then a CMOS process will simply be a silicon process with some of the bipolar transistor steps left out. But if resistors are needed, they can be added easily. And if a mostly ECL chip needs some MOS transistors, they can be added at little extra cost.

This should happen in the fairly near future, so it may be pointless to debate whether the future high-speed RAM technology will be ECL or biCMOS, and whether CMOS will replace biCMOS or ECL as a logic technology. It is to be hoped that, with

this process integration, some standard will emerge for high-speed signal levels. Then it will no longer be necessary to deal with CMOS logic and dynamic RAMs, with their positive power supplies and signal swings, in the same system as ECL-type negative supplies and swings.

**GaAs MEMORY.** As to GaAs, that technology may eventually produce large stand-alone RAM chips. Already, small GaAs RAMs have found some use in supercomputers and even in workstation servers. But ECL and biCMOS are still the technologies of choice for fast static RAMs. Before that changes, two impediments must be overcome.

One is just the immaturity of high-transistor-count GaAs technology. With smaller GaAs wafer sizes, larger cell sizes, and lower yields than an equivalent silicon wafer, it is not yet cost-effective to produce GaAs RAMs in higher densities. But GaAs processes are becoming more refined; they are likely to produce RAM cell geometries close to those of the most compact CMOS cells in the near future.

The second impediment will take some industry cooperation to overcome: incompatible signal levels. Commercial high-density GaAs logic chips handle external signals at ECL and TTL levels; but they vary widely from product to product and vendor to vendor in internal signal levels (anywhere from greater than zero to near -5 V). I/O level translation causes delay that robs GaAs of part of its speed advantage, and this will continue unless a standard direct GaAs signal level convention is developed.

Meanwhile, GaAs logic devices should continue to advance rapidly in gate count. GaAs gate density started out low and chips

started out small, but for good reason. Their developers had enough trouble building good GaAs MESFETs of any size and could not press the state of the art in lithography or expect a large defect-free chip area. (MESFETs, or metal-semiconductor FETs, are junction devices, in contrast to MOSFETs, or metal-oxide-semiconductor FETs, in which an insulating layer separates the metal from the semiconductor.)

But there is nothing that would prevent GaAs MESFETs from attaining the same small geometry as silicon MOSFETs. Already, individual MESFETs are nearly as small as MOS transistors. What GaAs needs is silicon-like contact and interconnection methods that let more transistors be hooked up in a given area. Extras like the GaAs equivalent of polysilicon resistors would also help RAM cells. And, of course, larger, more economical GaAs wafers are on the way; GaAs wafers are now about 10 cm in diameter, and they are likely to increase to 15 cm next year—smaller than today's 20-cm silicon wafers, but gaining.

When it comes to NAND gates and similar series circuits, GaAs junction FETs need extra diodes, whereas silicon MOSFETs do not. Here, silicon-based technologies may retain some superiority in functional density, but not nearly as much as in the past.

**TOP OF THE LINE.** CML GaAs, now neglected, may prove viable in the future for the highest-performance supercomputers. In the recent past, the density of CML GaAs chips has been too low; too many CML chips are needed, and the resulting interconnection delay offsets the logic's tremendous speed—already about 1 GHz. But as DCFL GaAs approaches biCMOS densities, CML



## A dotted kind of logic

For its next generation of supercomputers, to be introduced in about four years, Cray Research Inc. will use a new form of silicon bipolar IC. The emitter-coupled logic (ECL) so widely favored for high-speed computation will give way to emitter and collector dotted logic (ECDL). "Dotted" refers to the fact that several emitters or collectors are physically connected to form a logic function.

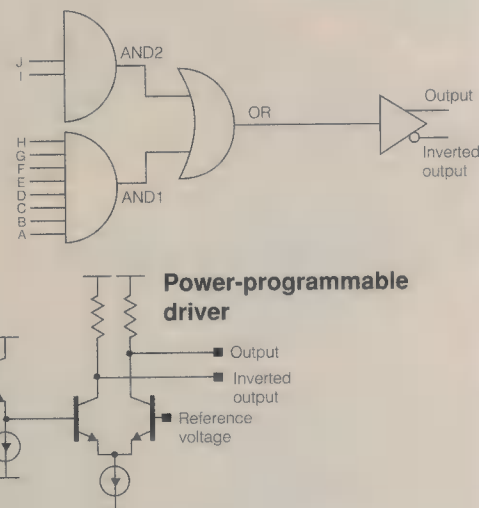

ECDL is both faster and more efficient than ECL. Clock rates above 500 MHz (that is, clock cycles less than 2 ns) have already been demonstrated.

An ECDL cell performs two levels of logic: emitter-dotted wire-AND gates are connected to a collector-dotted wire-OR gate [see figure]. As the signal gain in the logic circuit is less than 1, each cell is provided with an amplifying current-switch driver. In ECDL, the signal is restored in the amplifier after two dotted logic stages; in ECL, in contrast, each stage contains a signal-amplifying current switch.

The logic and driver circuits operate with low-swing signals of 500 mV on a low-voltage supply of -2.7 V. These low levels, and the absence of ECL-type series gating, account for ECDL's ultrashort gate delays and low power dissipation (65-90 ps and 0.7-0.8 mW per gate). Comparable ECL gates, fabricated with the same silicon bipolar process and subject to identical loading, experience gate delays of 100-150 ps at 2-3 mW per gate.

All switching within an ECDL circuit is in phase; that is, an output is high when its input is high. As soon as an input begins changing, the internal nodes begin changing in the same direction. The internal signal is not inverted as in ECL, and there is no threshold that must be passed before the internal signal begins changing. This in-phase characteristic gives ECDL its high fan-in capacity. It is not subject to the delay that occurs in inverting nodes when multiple inputs switch simultaneously. Moreover, ECDL input transistors are connected in an emitter-follower configuration that presents a high input impedance and minimum load to the preceding driver.

### Emitter collector dotted logic cell



Emitter and collector dotted logic (ECDL) creates a sum of products in negative logic. For the 8+2 input cell shown here, first-level AND gates are formed by linking emitters to nodes AND1 and AND2. A second-level OR gate is formed by linking diode-connected collectors to the node labeled OR. The power-programmable driver produces both true and complementary output signals for distribution on the chip.

Only one transistor is needed for each input, and all input transistors for a logic cell may be placed within the same collector region, or tub. This leads to compact logic layouts on the chip. And the regular pattern of the input transistors greatly simplifies generation of macrocells.

Each cell's current-switch driver may be set to any of eight programmable current levels to minimize the power consumption and thereby achieve the required delay for the interconnection-metal load and fan-out of the cell. The driver is programmed by computer-aided design software after the logic macros have been placed and the cell interconnections have been routed. The power programming keeps loading delays as low as 18 ps/mm of interconnect

and 1.1 ps per fan-out load at a 3.2-mA switch current.

The current-switch driver produces output signals having symmetrical rise and fall profiles. The symmetry decreases transient noise induced on power buses and adjacent signal lines to one-twentieth that of an emitter follower.

We are now making arrays of ECDL containing more than 50 000 logic gates and 300 input and output drivers on a 10-by-10-mm silicon chip, for a total of more than 1 million devices (resistors and transistors). Also built into the chip are a fixed clock distribution network and a boundary scan register for testing.

—Jan Wikstrom and Tony Vacca,  
Cray Research Inc., Chippewa Falls, Wis.

GaAs should draw closer to ECL logic density. When that happens, CML's speed may well compensate for its high power dissipation in top-of-the-line supercomputers.

Nevertheless, it is impossible to speculate on the longevity of silicon versus GaAs. Both materials have their pros and cons. The higher carrier mobility of GaAs lets it switch more current with smaller devices. Silicon can grow an insulator for the gate dielectric of insulated-gate FETs and its bipolar transistors are backed by years of evolution.

A trend that may favor silicon logic in supercomputers is the growing popularity of massively parallel processing (MPP), in which large numbers of small, lower-performance processors operate in parallel, replacing a single large processor. A good, cost-effective choice for CPUs in an MPP supercomputer may be the same reduced-instruction-set computer (RISC) chips used in workstations—relatively inexpensive silicon CMOS chips today and biCMOS in the next generation.

More certainty is possible over the role of customized versus standard memory in superconductor CPUs. As suppliers fit more and more transistors into the chip, much high-speed RAM has moved onto the logic chips, and this trend seems likely to continue. Any discrete memory chips will take the form of application-specific RAMs—that is, RAMs designed for a particular application in a particular supercomputer. Supercomputer manufacturers are forced to optimize RAMs for individual applications instead of making do with less-than-optimum catalog parts. Today, IC suppliers are working directly with those manufacturers on custom RAM designs, instead of offering new, very high-speed, standard RAMs to the market.

Will wafer-scale integration show up in supercomputers? The technology was widely touted a few years ago. The principle behind it is to put a complete subsystem on a single semiconductor wafer by processing several different device types on it and by

interconnecting the good devices as required. So far, however, the technology has not been cost-effective. Just one good device too few of a required type, and the whole wafer has to be discarded. Enough good devices of all types, and the wafer is likely to contain unused good devices and have a low functional density.


If wafer-scale integration has a place, it may be in some mass-memory applications, but not in the CPU of a supercomputer. Multichip modules (MCMs) are a far better alternative. MCMs combine several bare chips in a single package where chip-to-chip interconnections enjoy a favorable environment designed to minimize signal delay.

**ABOUT THE AUTHOR.** Harold Dozier is vice president of advanced development at Convex Computer Corp., Richardson, Texas. From 1988 through 1990, he managed the design of the GaAs-based C3800 supercomputer. He is now selecting an application-specific IC technology for a future massively parallel processing product. ♦



# Interconnections and packaging

*Designers are rushing to remove interchip and intermodule bottlenecks as I/O links grow in number and heat loads soar*



The 1-gigahertz clock rates all but assured for high-performance computers by the next century will set a fast pace for packaging engineers, who must somehow transfer signals between chips at rates compatible with those blinding internal speeds. To avoid or minimize chip-to-chip delays, IC manufacturers are packing ever greater functionality on chips, and supercomputer manufacturers are packing chips closer together on circuit boards.

The advanced interconnection, packaging, and cooling technologies now being developed for supercomputers have profound implications for computers in general; many innovations in supercomputers are likely to trickle down to minisupercomputers, mainframes, minicomputers, workstations, and even personal computers. Examining hardware trends in supercomputers is like gazing into a crystal ball at tomorrow's ordinary computer.

To the packaging engineer, the demands of supercomputers mean that many more input/output connections must be squeezed into less space and that greater quantities of heat must be removed—again, in less space. “Getting the heat out of high-performance devices while maintaining a temperature rise of less than 50°C is an enormous challenge,” Lance Glasser, director of the Electronic Systems Technology Of-

fice, Defense Advanced Research Projects Agency (Darpa), Arlington, Va., told *IEEE Spectrum*. “At 20 W/cm<sup>2</sup>, a supercomputer IC generates six times the heat flux of a 60-W light bulb.” Many see that heat flux increasing to 100 W/cm<sup>2</sup> within five years.

**EXPLORING.** To cope, engineers are adding more interconnection layers. They are turning to multichip modules (MCMs). They are adopting exotic materials such as diamond, and they are exploring new technologies such as high-temperature superconductivity.

“We’ve hit a wall of latency,” Burton Smith, chairman, Tera Computer Co., Seattle, Wash., told *Spectrum*. Latency, a measure of the tendency of interconnections to delay high-bit-rate signals, prevents advanced microprocessors from running as fast as they could. “Today, no one knows what to do with a 400-MHz clock because of latency in interconnections,” Smith said. “Gallium arsenide IC manufacturers are really targeting their products for low power at silicon ECL [emitter-coupled logic] speed; that is, they’re not trying to go faster than the competition.”

To lower latency, supercomputer manufacturers are making circuit boards out of new polymers and ceramics with lower dielectric constants. Benzocyclobutenes (BCBs), for example, have only half the dielectric constant of the usual polyimide-glass or epoxy-glass board, absorb far less moisture, and may be etched by plasma as well as chemicals. BCBs withstand temperatures above 430 °C, a big help with high-heat-dissipation devices. Incorporating ceramic fiber fillers like aluminum nitride, silicon carbide, and beryllia in such polymers will raise their thermal conductivity and reduce their thermal expansion. BCBs are a candidate, too, for on-chip thin-film insulation between metallization layers, as are polyimides with low expansivity, low stress, and low dielectric constant.

Another promising board material is synthetic diamond, which is an electrical insu-


lator and thermal conductor. Diamond is four times as good as copper at conducting heat. Darpa is sponsoring research on diamond manufacturing at Norton Co., Worcester, Mass., and General Electric Co., Schenectady, N.Y. The goal is to produce affordable diamond substrates for three-dimensional supercomputer modules with high density and high clock rates.

“By 1993, the program hopes to achieve thick 10-cm-by-10-cm substrates at less than [US] \$200/cm<sup>2</sup>,” William Barker, who runs Darpa’s diamond program, told *Spectrum*. “At slightly lower cost, diamond would be good for thermal management in high-performance workstations.” The eventual goal is substrates costing less than \$2/cm<sup>2</sup>. The chief constraint is the cost of energy for chemical vapor deposition; the raw material—carbon—is cheap.

Darpa is now considering applications and design concepts as well as techniques for metallization of diamond. The agency is also studying methods for forming multiple layers of interconnections and through-the-substrate vertical interconnections in 3-D diamond multichip modules.


Three-dimensional MCMs take advantage of the thinness of IC chips to provide short, vertical chip-to-chip connections. According to this scheme, arrays of chips, each about 0.6 mm thick, would be piled one on another. For example, a chip that in a two-dimensional array, has eight near neighbors, could have 116 chips within the same interconnection distance in a three-dimensional array.

Richard C. Eden, a consultant to Darpa based in Thousand Oaks, Calif., gave an example of the effectiveness of an extra dimension in reducing interconnection delays: a three-dimensional array of 117 chips, each containing 10 000 gates, would place more than 1 million gates within 650 ps of each other if the interconnection delay were about 60 ps/cm, a representative value. This proximity could readily support a clock rate of almost 1 GHz.



[1] Vitesse Semiconductor Corp.'s plastic package for GaAs gate arrays comes in surface-mount [shown here] and pin-array versions.





[2] Multichip module, built as a demonstration of quality and yield, rather than density, contains 10 meters of 10- $\mu$ m-wide high-temperature superconductor interconnections. It was made by Superconductor Technologies for the Defense Advanced Research Projects Agency.

key component is a molded Kevlar-filled polyethersulfone chip carrier, seated on a strip-line printed-wiring substrate with nine layers. The assembled package is about 5 cm square.

Another way to lower latency is to use

thinner, shorter conductive paths between chips. Thinner conductors also allow denser interconnections, and hence higher chip pinouts. Using photolithography, circuit board fabricators routinely turn out 0.05-mm-wide conductors and stack them 12 layers high. Unfortunately, the thinner conductor, the higher its resistance, the lower its current capacity, and the greater its sensitivity to noise.

That is why supercomputer manufacturers are looking to enhance high-density MCMs

On the other hand, if the same number of chips had to be spread out two-dimensionally, the longest delay would be more than 1.7 ns. That delay would preclude operation at rates much above 250 MHz.

**CUBE COMPUTERS.** Three-dimensional MCMs will lead to complete supercomputers and workstations in a compact cube, perhaps 10 cm on a side, Eden believes, but the modules must be interconnected vertically in such a way that they can be removed and replaced. Equally important, the intense heat that the chips generate has to be removed efficiently to the faces of the cube, where it can be carried away; there will be little space within the densely packed cube for coolant to flow.

"This is where diamond MCM substrates can help," Eden told us. "Diamond laterally conducts the heat generated by the logic out of the actively populated cube." With diamond, there is no need for coolant channels through the cube; the full cube volume can be utilized for interconnections. The packaging designer can use the simplest, most reliable, least expensive demountable interconnections available without being constrained by the need for routing coolant.

**BACK TO PLASTIC.** Low-dielectric-constant materials are important in supercomputer device packages as well as in interconnection boards. Vitesse Semiconductor Corp., Camarillo, Calif., chose plastic for a new GaAs gate-array package [Fig. 1] because the material has a lower dielectric constant than the ceramics used earlier for high-pinout chips (the plastic package is cheaper to manufacture, too).


Vitesse's 557-pin package accommodates signals at rates up to 600 MHz. It holds a 350 000-gate GaAs chip dissipating 44 W. A

to help with their chip-to-chip communication problems. By packaging anywhere from two to over 100 bare chips in a controlled electrical and thermal environment, MCMs reduce delays between chips, simplify power distribution, and help to dissipate more heat. Darpa regards MCM technology as one of its top priorities. Although only a small fraction of chips are mounted in MCMs at present, Darpa expects the share to grow to 30 percent by the year 2000 and even larger in the first decade of the new century.

**NO CAPACITORS.** To further exploit these advantages now, developers are studying new MCM structures and materials. At Rensselaer Polytechnic Institute's Center for Integrated Electronics, Troy, N.Y., researchers are evaluating, under Darpa sponsorship, ferroelectric ceramic films as replacements for space-consuming capacitors, now needed to decouple signals. The film would take up no MCM area. The center is also researching novel substrate structures with low dielectric constants for IBM Corp.

Of course, narrow interconnections in MCMs create added resistance, just as they do on circuit boards. But if the interconnections are superconducting, they have no resistance and may be made much thinner. Two layers of high-temperature superconducting interconnects can replace dozens of ordinary conductor layers. Superconducting MCMs are therefore potentially denser and easier to manufacture.

**CHILLING CMOS.** By next year, Darpa hopes to demonstrate that dozens of silicon CMOS chips can be interconnected by superconduc-



[3] Paragon massively parallel-processing supercomputer needs hardly any backplane cables; ICs route signals among processors.



tors cooled by liquid nitrogen. Chilling CMOS to 77 K has another plus: it increases switching speed, especially when the CMOS channel lengths are much less than 1  $\mu\text{m}$ . "If a computer could run at a 100-MHz clock rate normally, then 250-MHz should be achievable using high-temperature superconductor MCM packaging," said consultant Eden.

Superconductor Technologies Inc., Santa Barbara, Calif., is developing thallium-based superconducting MCMs for Darpa. So far, the company has achieved 10- $\mu\text{m}$  line widths [Fig. 2]. It is also developing, with Sunpower Inc., Athens, Ohio, a long-life, low-cost, compact cryorefrigerator. Based on a closed-loop helium Stirling cycle, the cooler is about 8.5 cm square and 26 cm long and has a 3-W cooling capacity. James Long, Superconductor Technologies' vice president of contract marketing, foresees an increase in capacity to 5 W and commercial applications in five to seven years. The primary market will be supercomputers, but high-end workstations will also use cryocooling, Long believes.

Another coolant candidate is liquid nitrogen in an open cycle. Nitrogen in the air would be liquefied, pumped to the MCMs, and allowed simply to evaporate back into the atmosphere. Tony Vacca, vice president of technology at Cray Research Inc., Chippewa Falls, Wis., thinks this may be a viable cooling technique for small single-processor applications. But, "for extremely large supercomputers, a lot of work will have to be done because we're talking over 100 000 W, and that's a lot of nitrogen," Vacca said. "One can picture a computer site with a large storage tank in the back if the system is in the open-loop mode. Or one can picture a high-power-consumption refrigeration unit if the closed mode of operation is used."

Whether cryogenic or not, cooling will remain essential to supercomputers, although coolant materials are likely to change. Fluorocarbons unfriendly to the environment are likely to give way to more benign cooling media. "We require fluids that don't attack the environment," Vacca said. "We've moved into that arena already. One fortunate thing is that water has never been declared environmentally dangerous; you can pull a lot of heat away with it."

Devising adequate interconnections for ever-rising levels of chip integration will keep packaging engineers busy, although they know generally what to expect. Paul A. Totta, an IBM fellow at East Fishkill, N.Y., predicted that, if current trends continue during the decade ahead, the highest levels of integration in silicon chips will be gigabit dynamic RAMs, 10-million-gate CMOS logic, and 100 000-gate bipolar chips. "There will also be equivalently large microprocessors, biCMOS chips, and all kinds of custom, hybrid mixtures of logic and memory," Totta said.

Time-honored wire bonding will remain popular for carrying signals to main memory chips, Totta forecast, because of its low cost and high reliability—and memory chips' relatively modest I/O requirements. But "area-array flip chips will become more widespread on multichip substrates because of the high I/O requirements of ULSI [ultralarge-scale integration] logic chips—3000 to 5000 pads—and the better properties of the short electrical path," he told us. **USING TOTAL AREA.** In area-array flip chips, the entire IC area, not just its periphery, is used for on/off chip connections. Small sol-

## Interconnection delays mean that today no one knows what to do with a 400-MHz clock

der bumps are attached to pads on the chip surface; they may even be placed over active devices. The chip is inverted and the bumps are bonded to interconnections on an MCM substrate. IBM has long advocated area bonding, and other companies are beginning to see its benefits. "Area bonding will become dominant because of the enormous number of gates we can put on future semiconductor devices," Cray Research's Vacca told us. "You just run out of pins with edge bonding, but the area of the die provides a new opportunity. So a lot of R&D is going into using the total area of the chip."

**STAYING 2-D.** Intel Corp.'s Supercomputer Systems Division, Beaverton, Ore., however, eschews exotic technologies in its massively parallel-processing (MPP) computers. "In terms of the future, we're trying to stay with interconnection networks that are easy to embed in the backplane," Justin R. Rattner, the division's director of technology, said. "We're trying to avoid 3-D packaging arrangements because they inevitably result in a large number of cables." Cables raise costs, reduce reliability, and limit performance, in Rattner's view.

A case in point is Intel's new Paragon MPP machine. Instead of the usual backplane with hundreds of copper cables and traces across the boards, Intel uses an active backplane with very large-scale integration (VLSI) chips that route and relay communications traffic between processors [Fig. 3].

Each 325-pin IC simultaneously routes information in two dimensions and provides a high-speed connection to a single processing node. The chips were developed and manufactured by Intel from designs originated at California Institute of Technology, Pasadena, by Charles Seitz and his graduate students.

The unidirectional physical links between the routing ICs and between active backplanes are either short circuit-board-like traces or point-to-point flexible circuits. There are no buses in the conventional sense. "Because we're dealing with short, direct connections, we're able to run the wires much faster than one would normally expect for the chip technology we use," Rattner told us. The longest flex circuit jumper extends no more than 15 cm. "We're using 50-MHz chips and running the wires about twice that speed."

Active backplanes yield an uncluttered equipment cabinet, remarkably free of the tangle of interconnection cables that characterizes most supercomputers. "When we open the back door on a Paragon machine for people who have seen a traditional supercomputer, their jaws just drop," Rattner told *IEEE Spectrum*.

Intel's philosophy is to use technology adaptable to desktop machines wherever possible. Processing nodes on the Paragon resemble motherboards on PCs, for example. The only distinguishing feature is the presence of a

high-speed interprocessor communications interface. The packaging and materials are conventional. The rationale is that high-performance computers will have similar architectures, and the competitive advantage will go to the manufacturer that has the lowest-cost technology and will come from high-volume applications at the desktop level.

Meanwhile, work proceeds on optical-fiber interconnections which plug into every level. "We see that technology playing a very big role very soon in box-to-box interconnection—between stand-alone units 10 meters apart or 20 km apart," Daniel Stigliani told us. Stigliani, a senior technical staff member at IBM's Enterprise Systems, Poughkeepsie, N.Y., foresees optical fibers becoming an interconnection medium on another level—between circuit cards and between MCMs—before the end of the decade. "It's clear that data processing will require transmission at gigahertz and beyond," he said. "The only technology that can satisfy these needs is fiber-optic technology."

Much remains to be done. Board-to-board and MCM-to-MCM fiber links will need less costly electro-optic conversion devices, as well as inexpensive, mass-producible optical alignment techniques. And routing fiber is no picnic. "You don't bend fibers at the radii you do copper," Stigliani said. "But everybody who's interested in high-speed data communications has by necessity to get into this."

Darpa, again, clearly agrees. The agency has formed a consortium to research optical-fiber interconnections over the next 30 months for imaging systems and parallel processors, as well as for telecommunications switching. The Optoelectronic Technology Consortium's members are AT&T, General Electric, Honeywell, and IBM. ♦



# Beyond today's supercomputers

*Optical input/output and chips grown from organic molecules are likely—or a machine may exist only when called into being*



Technologies so fantastic they sound like science fiction may mold the next generation of supercomputers. A poll of researchers investigating the frontiers of this new world reveals bacteria serving as lithographic masks, computer chips grown from organic molecules, beams of

light in free space acting as input/output "buses," and architectures without logic gates that yield an answer without calculations.

A surprising number of people in the field believe that supercomputers as they are today—single, multiuser machines in one room—may disappear. Instead, users may "construct" a supercomputer ■ needed at a moment's notice; they would get ■ national or international multigigabit-per-second optical-fiber network to call on many computers for coordinated number-crunching. Alternatively, the power of a supercomputer might be contained in ■ single desktop workstation.

According to the investigators, some of the technologies may bear fruit in high-performance experimental machines within the next 10–15 years.

**LET'S GET SMALL.** In today's supercomputers, the speed of signals on chip and elsewhere is already approaching the speed of light. Therefore, the only way to further increase computer speed is to shrink device size so as to minimize the signals' "commuting distance" and hence their travel time. Memories also

need to grow smaller and denser: the goal for future supercomputers is the terabit memory—one with more than a trillion ( $10^{24}$ ) bits stored in ■ square centimeter.

But at feature sizes of 10–100 nm, the resolution of even X-ray and electron-beam lithographic techniques becomes too coarse. "You start seeing deviations in the edges of the masks on the scale of an electron's wave-

length," said Mark Reed, professor of electrical engineering at Yale University in New Haven, Conn. Reed and others are studying how to make masks with nanometer-scale patterns that give sharp-edged images.

One way, pioneered in Austria, is to fabricate masks using the two-dimensional crystalline layers that encase many bacteria. Uwe B. Sleytr and his colleagues are following this track at the Center for Ultrastructure Research at the University of Agriculture in Vienna. Such surface-layer lattices are assemblies of a single protein or glycoprotein species and show oblique, square, or hexagonal symmetry with center-to-center spacings of 5–32 nm.

Moreover, when the surface layers are removed from the bacteria and held in suspension or laid flat on surfaces to which they can adhere, they neatly assemble themselves into immobile two-dimensional arrays of identical subunits. Reporting at the Second International Conference on Molecular Electronics, held on St. Thomas in the U.S. Virgin Islands, Dec. 15–19, 1991, Sleytr and company observed that the bacterial surface layers "are nature-tailored patterning elements for nanometer technologies."

**GLASS SPONGES.** Another technique for making nanometer-scale masks is inspired by optical fibers, which are formed by pulling heated glass rods into the thinnest of threads. For the masks, ■ boron-rich glass core is surrounded by a leaded-glass cladding. The

nels smaller than 33 nm in diameter.

When exposed to an acid, the etchable glass melts away, leaving a honeycomb of empty channels. This nanochannel matrix can be cleaved into slices ■ few nanometers thick, affixed flat on a substrate, and used as ■ mask in the fabrication of large arrays. Depending on the materials used, the glasses are stable at temperatures over 700 °C. Since that is higher than typical semiconductor processing temperatures, the glass masks have a distinct potential for growing nanometer-scale structures using such atomic layering techniques ■ molecular beam epitaxy, chemical vapor deposition, and ion implantation. Moreover, the nanochannel glass could provide ■ packing density approaching that necessary for a terabit memory in ■ future supercomputer.

Another use for the nanochannel matrix is to have a crystalline semiconductor, such as gallium arsenide, deposited in those channels to produce ■ regular array of millions of quantum wires or quantum dots. (Quantum wells, wires, and dots are structures useful for trapping and quantum-mechanically confining electrons or other carriers to produce certain effects not observed in bulk semiconductor materials. A quantum well is ■ two-dimensional layer, ■ quantum wire is a one-dimensional line, and ■ quantum dot is a zero-dimensional point. Large arrays of quantum dots are envisioned as ■ key to high-performance computing, either

■ switches or as dense memories.)

Moreover, since the glass holding those structures is transparent to visible and near-infrared radiation, the nanochannel matrix is promising for ultradense optical and electro-optical quantum-confined devices including optical memory and optical I/O, Tonucci added.

**GROWING IT IS EASIER.** But lithography of any kind may be the hard way to obtain the nanometer-scale structures necessary for a terabit memory of a future supercomputer. On the terabit

scale, "the total surface area for the storage of a bit of information starts being the size of ■ large molecule," observed George Whitesides, professor of chemistry at Harvard University in Cambridge, Mass. "So, if you want to take a terabit memory seriously, you are almost driven to look at quantum devices where the basic storage device is a molecule, and you put an electron in or take one out" to change the state.

That scale of device cannot be built eco-

*'Many people believe the days of the supercomputer as a large...shared machine are numbered'*

core is etchable with an acid, the cladding is more inert, said Ronald Tonucci, research physicist at the Naval Research Laboratory in Washington, D.C.

First the two-part glass is pulled to the thinness of ■ hair, is cut into lengths, and the lengths bundled together. The bundle is then pulled to the same thinness, cut, and bundled, repeatedly 5–15 times. At that point, Tonucci noted, the result is a matrix of inert glass holding up to 10 million uniform chan-





Purdue University

A research prototype of a parallel supercomputer capable of choosing which type of parallelism best suits each part of a problem has been devised by H.J. Siegel [far right] with his colleagues at the Parallel Processing Laboratory at Purdue University. Known as PASM-1 (for Partitionable SIMD/MIMD), the prototype can be dynamically reconfigured to operate as one or more independent SIMD (single-instruction-multiple-data streams) and/or MIMD (multiple-instruction-multiple-data streams) submachines of various sizes. Each PASM submachine can switch independently between the two modes of parallelism during program execution, an ability opening the new research area of mixed-mode computation. PASM-1 consists of 30 processors, but the concept can support up to 1024 processing elements.

nomically with conventional methods of fabrication. "Much of the expense in attaining small feature sizes is the cost of elaborate clean rooms and complicated processes," Whitesides said. One particle of dust or even a vapor can be catastrophic, especially to the regular arrays of large memories, he said.

Whitesides is looking at chemical processes in which defect-free molecular structures assemble spontaneously because that is their nature. A crystal, such as silicon or gallium arsenide, is the simplest form. But far more varied in their structure, said Whitesides, are the vast range of organic molecules: molecules containing carbon and compounds found in living organisms.

Organic materials have a number of attractive properties. First, the structures they form "intrinsically reject contamination and heal themselves of defects," he said, because chemical processes are highly selective. So high-grade clean rooms may not be necessary for their fabrication, which would be done mostly in chemical baths.

Moreover, many organic molecules have excellent electronic properties, Whitesides said, "especially if you talk about energy transfer (moving an electronic excitation from place to place) instead of electron transfer." He also envisions hybrid systems, such as colloid particles of gallium arsenide coated with an organic material that self-assembles into a flat array of quantum dots.

The surface layers of bacteria or the crystalline form of a virus are "actually too big for a terabit memory," Whitesides said. The appropriate size is a protein.

Along these lines, Tadashi Ishibashi, of Tokyo-based Hitachi Ltd. is investigating the use of rhodopsin as a fast optical switch. Rhodopsin is the light-harvesting protein of certain photosynthetic purple bacteria (those that turn salt ponds red); it is also the "visual purple" in the retina of human and animal eyes allowing night vision. Ishibashi's project is sponsored by Japan's R&D Association for Future Electron Devices as part of the Research and Development Project of Basic Technology for Future Industries, initiated by Japan's Ministry of International Trade and Industry in 1981. Ishibashi and his colleagues are developing a method for the two-dimensional ordering of rhodopsin molecules on a substrate, using an antibody to immobilize the photosensitive rhodopsin on a lipid film.

"We can make self-assembled structures, and we can make electronically functional organic molecules. The next step is to make self-assembled electronically functional structures," said Whitesides. Another challenge is soldering leads to biochips for reliable electric contact, added Yale's Reed.


**NO MORE WIRES.** Most researchers agree that the long, high-bandwidth interconnections (for communications, say, from processor to processor or over local-area networks) in a future supercomputer will have to be optical. Optical methods offer more bandwidth, less noise, and greater electronic isolation among connectors than do metal wires, explained Harold Stone, research staff member at IBM Corp.'s Thomas J. Watson Research Center in Yorktown Heights, N.Y.

"Optical fiber has a bandwidth of about 35 000 Gb," Stone observed. "We use only a gigabit of that when we have a gigabit network" such as the high-speed data highway proposed by Senator Albert Gore (D-Tenn.).

According to Stone, today's supercomputers are pressing the limits of the bandwidth and clock cycles that metal wires can convey outside a chip, because of noise from three sources: transients from the turning on and off of devices, reflections from the terminations of circuits, and cross talk among the electric fields of tightly packed wires.

The noise worsens with clock frequency, interconnection density, and interconnection length, to the point that at clock cycles of a few nanoseconds, it can prevent reliable transmission. Clock rates are now approaching frequencies where these noise sources matter to metal wires as short as 10 cm. Moreover, wires thinner than 0.1  $\mu\text{m}$  begin to show quantum effects, behaving more like waveguides to individual electron wave functions than as a conduit for current.

Optical waveguides do not suffer from those problems, and thus offer higher density and the potential to be clocked at a higher rate than metal interconnections. Stone's focus therefore is on replacing metal wires with optical waveguides.



A nanochannel glass structure has been developed at the Naval Research Laboratory in Washington, D.C. The two-dimensional array consists of glass capillaries in cross section, formed by repeatedly stretching glass fibers whose core of etchable glass was later removed by acid. In this scanning electron micrograph, made by the lab's Ronald Tonucci, each capillary is 33  $\mu\text{m}$  in diameter. For future supercomputers, these regular arrays could be used for masks in semiconductor processing, for a terabit memory, or for ultradense optical and electro-optical quantum-confined devices.



Not even fibers may be necessary for input and output, said John Caulfield, university eminent scholar at Alabama A&M University in Normal. Among other things, he is interested in a free-space optical method for addressing the individual organic molecules that store information in a terabit memory of the kind envisioned by Whitesides. Caulfield is investigating this technique in cooperation with Urs Wild at the Swiss Federal Institute in Zurich, Edward Manykin of the Russian Science Center (the I.V. Kurchatov Institute) in Moscow, and Valentin Morozov at the University of Colorado in Boulder.


"Even a 1- $\mu$ m-diameter beam is a blunt instrument" compared with a protein molecule, Caulfield said. It would shine on thousands of molecules. "But what if you viewed it not as one probe, but as many probes?" Caulfield asked. The beam of a tunable dye laser "could write things at many wavelengths simultaneously."

If each molecule within the beam responded to only one combination of wavelength and other parameters, information could be stored on a surface of many molecules, "but any piece is approaching the size of a molecule," Caulfield said. An analog is a radio receiver, tuning into only one channel out of the many it can receive.

Optics also holds out hope for quantum-mechanical computing: the ability, for special problems, to do away with logic gates altogether. The trick is to choose "a physical process that nature knows how to carry on, and make it correspond to a mathematical operation," Caulfield said. One such example in analog computation is the Fourier transform: the resolving of a function into its sine-wave components by directing light from a laser through an image and then through a lens; the focal plane of that lens will contain a spatial display of the Fourier transform of the image, needing no numerical calculation. In such specialized cases, "you set up the computer to solve the problem optically, then you insert the data for the particular problem," Caulfield said. "The measurement of the intensity of the light at the output plane is the solution."

In digital computing, another example is a programmable logic array. According to Caulfield, the key operation is a massively parallel OR gate, which asks one question: did light reach this detector from any of a million inputs? To electronics, that is a 1-million-part question, Caulfield observed. But to optics, a yes or no, given by the detection of light or no light, can be answered with only 100 photons with a 1-part-per-billion accuracy, he said. "This is quantum mechanics' two-slit experiment extended to a million slits," said Caulfield, as each photon "knows about" all 1 million slits.

The price? "You give up asking certain questions," he admitted; in doing an opti-



*Molecules of rhodopsin, a photosensitive protein found in retina cell membranes, could be the basis for biomolecular devices in possible future supercomputers, according to Tadashi Ishibashi and his colleagues at Hitachi Ltd. in Japan. When light strikes rhodopsin, chemical changes generate an electric signal. To retain its photoreactivity outside the eye, rhodopsin must be held in a specific orientation when bound to a substrate. In a project backed by Japan's R&D Association for Future Electron Devices, Ishibashi's group took molecules of rhodopsin from the eye of an octopus and immobilized them on a two-dimensional lipid monolayer with a hybrid antibody, which binds preferentially at one end to the rhodopsin and at the other end to the lipid. (Source: Future Electron Devices Journal, Vol. 2, Supplement, 1992, p. 53.)*

cal Fourier transform, say, the answers for intermediate computations are lost, and in some circumstances they might be useful. "But if you're not too curious about what happened and only care about the final result," the payoff is an answer literally at the speed of light—about 1 ns in an apparatus

**'It is difficult and probably a bad idea to build conventional logic on the nanometer scale'**

3 cm long—plus a complete avoidance of building devices and cooling them for the intermediate computations.

**NOWHERE AND EVERYWHERE.** "Many people believe the days of the supercomputer as a large multiuser shared machine are just as numbered as the day of the mainframe," said Robert Hecht-Nielsen, chair of the board of Hecht-Nielsen Neurocomputer Corp. (HNC) in San Diego, Calif.

Others concurred. "It's hard to distinguish a supercomputer from a workstation these days except in quantity of processors, not speed," agreed IBM's Stone. He pointed out that the Cray Y-MP, a multiprocessor

supercomputer released in 1988, has a clock cycle of 6 ns, while the Alpha microprocessor for workstations and massively parallel processors, which Digital Equipment Corp. of Maynard, Mass., announced this year, has a clock cycle of 5 ns. In fact, Cray Research Inc. in Eagan, Minn., the world's largest supercomputer maker, has publicly stated that its first massively parallel processor will be based on arrays of the Alpha chips.

Alternatively, "the supercomputer of the future is going to be distributed all over the world," declared Alan Huang, head of the optical computing research department at AT&T Bell Laboratories in Holmdel, N.J.

"What's going to make that possible is optical-fiber networks: intellectual power lines. CPU [central processing unit] clock cycles will be traded the same way as the power utility grids now trade kilowatts," Huang predicted. "You log onto the network and you're not going to know where the functions are done—your CPU may be in Kansas and your data on a disk in Saskatchewan," just as someone in New York State may be using power generated in Canada. "The only reason people may not think of it as a supercomputer is that it's spread across a continent," he said.

In Huang's view, the sheer power for distributed supercomputing power exists even now in a million processors worldwide: the main challenge is hooking them together. Part of that challenge is the time lag (or latency) of the signal traveling thousands of kilometers at the speed of light, and how latency affects the simultaneous scheduling of many tasks. Other challenges include mechanisms for billing for the time.

For handling the scheduling, Huang is exploring whether to utilize the inevitable latency to create delay lines; then, using a technique he calls computational origami, the distributed supercomputer would "reshape" the algorithm to be used in computation to partition the problem efficiently onto a pipeline architecture.

Already, distributed supercomputing is being done experimentally using 5-Gb/s test-bed networks established over the last two years by the Corpo-

ration for National Research Initiatives, headquartered in Reston, Va. One network is the Casa Distributed Computing Environment. It links supercomputers at the California Institute of Technology, the Jet Propulsion Laboratory, the San Diego Supercomputer Center, and Los Alamos National Laboratories.

**WANTED: NEW MODEL.** Regardless of whether the future supercomputer sits on a desk or is a cross-continental virtual machine, one supreme desideratum remains: a universal computational model for parallel processing. For serial computation, the von Neumann model forms the conceptual basis for com-



putation, communication between the CPU and memory, and layout of the data. But the von Neumann model does not address certain needs of parallel computation, such as communication among many processors.

"Today, when you write a parallel program, there is no universal way you can specify your program to run on half a dozen different machines," said H.J. Siegel and Seth Abraham of Purdue University, West Lafayette, Ind. They are professors in the school of electrical engineering there, and Siegel is also coordinator of the school's parallel-processing laboratory.

Even more important, until such a computational model exists, supercomputers may not fulfill their greatest potential. The main reason is that efficient communication among the processors and the scheduling of tasks become crucial, as well as the type of parallelism and the number of processors.

Even absent a computational model for massively parallel systems, researchers are coming to realize that building structures on the nanometer scale yields quantum effects that demand entirely new architectures. "It is difficult and probably a bad idea to build conventional logic at 20 nm," declared Brosl Hasslacher, staff physicist in the theoretical division at Los Alamos National Laboratories in New Mexico.

Latching onto the fact that quantum dots communicate with their immediate neighbors by the tunneling of electrons from one to the next, Hasslacher has devised an architecture called a lattice gas. It consists of a hexagonal grid with a particle at each node. With each tick of a clock, all the particles leap at once in some random direction to another node on the grid. When two particles collide, at the next tick of the clock they carom off one another in directions specified by an input set of collision rules.

It so happens that this lattice gas, despite its simplicity, reproduces the full dynamics of an ideal fluid (including vortex phenomena, turbulence, and other instabilities) over a small range of fluid parameters, Hasslacher said. That makes it an extremely powerful architecture for modeling fluid dynamics, including flows in the ocean, atmosphere, and interstellar gases. "By accident, it's a natural architecture for massively parallel machines," he said, even though it does not rely on numbers. "It doesn't do floating point; you have to average over regions."


**PANEL OF EXPERTS.** "People use supercomputers to solve specific classes of problems," observed HNC's Hecht-Nielsen. "Many models are built from first principles: that is, the laws of physics and chemistry are used to build mathematical models to solve equations." But when the behavioral rules of a material or fluid under investigation are not well-known, working from first principles may not be the right approach, he said. Of more help would be a computer architecture that can observe phenomena and discern patterns or rules implicit in behavior or in large amounts of data.

A promising architecture for this purpose is a neural network, working with a number-crunching computer. A neural network parallel-processes floods of information. It consists of many processing elements, each with local memory. Each element is connected to many others, from which it receives stimuli—inputs and feedback—and to which it sends stimuli. Some of those connections are strong, others are weak. The topology and strength of the connections influence what information-processing functions the network can carry out.

A neural network is not programmed with step-by-step algorithms, but learns through supervised training. It is fed raw data, per-

Lee Giles, senior research scientist at NEC Research Institute Inc. in Princeton, N.J.

Already, optical associative memories are being used on ultrafast spatial light modulators to control non-von Neumann data-flow supercomputers in which the data triggers the instructions. The experimenters here are Alabama A&M's Caulfield plus a group headed by Vsevolod Burtsev of the Russian Academy of Sciences in Moscow. In the prototype (which Burtsev calls a commutator), simple optical neural network tests for the presence or absence of the needed data and when there is enough data to do a calculation, sends it to the next available processor, Caulfield explained. The computations



*This laboratory model of an optical interconnect for possible use in a future supercomputer was developed earlier this year by scientists at Physical Optics Corp. in Torrance, Calif., in cooperation with John Caulfield at Alabama A&M University. The interconnect itself consists of a planar waveguide; at one corner is imprinted a 1-by-1-mm hologram of dichromated gelatin. When the green beam of an argon laser enters the edge of the waveguide, it propagates as a planar traveling wave [wavy lines]. When the beam meets the hologram, it splits into 50 beams that are directed into the air perpendicular to the waveguide surface.*

haps through sensors, along with feedback on how well or badly it is doing. As it runs through the material again and again, it makes myriads of mistakes; but learning from them, it finally organizes itself to carry out the desired task on new data.

Neurocomputers are usually built by software simulations. They do well at some things conventional computers are bad at. They are good at recognizing complex patterns, identifying handwritten characters, monitoring complex processes, optimizing complex schedules, and determining that a target seen from different angles is in fact the same object. (Of course, conventional computers are better number-crunchers.)

In the future, neural networks may be used as special pattern-recognition components in high-performance computers for such applications as speech recognition, said

themselves, though, stay electronic. The result is an optically assisted electronic-multi-processor supercomputer.

Other investigators are experimenting with combinations or hybrids of high-performance machines. For example, Purdue's Siegel and his colleagues have built a small-scale 30-processor prototype machine called PASM (PARTitionable SIMD/MIMD), and are developing software for it that will partition a problem so that each stage gets the type of parallelism best suited to it.

**ACKNOWLEDGMENTS.** Thanks also go to Ravindra A. Athale, associate professor of electrical and computer engineering at George Mason University in Fairfax, Va., and Karl Hess, research professor in the University of Illinois's Beckman Institute for Advanced Science and Technology in Urbana. ♦



# To probe further

**GENERAL INTEREST.** The IEEE Computer Society and National Aeronautics and Space Administration will jointly sponsor "Frontiers '92: The Fourth Symposium on the Frontiers of Massively Parallel Processors," next Oct. 19-21 in McLean, Va. For details, call the Computer Society at 202-371-1013.

There is also an annual comprehensive conference on supercomputing, sponsored jointly by the IEEE and the Association for Computing Machinery (ACM). Copies of the proceedings from past conferences are available through the IEEE Computer Society Press, Customer Service Center, 10662 Los Vaqueros Circle, Box 3014, Los Alamitos, Calif. 90720-1264, or the ACM Order Department, Box 64145, Baltimore, Md. 21264; 800-342-6626. The next conference, Supercomputing '92, will be held Nov. 16-20 in Minneapolis, Minn.

The Winter 1992 issue of *Daedalus*, the journal of the American Academy of Arts and Sciences, was devoted to the future of supercomputing. For a copy, send US \$7.95 plus \$3 shipping to *Daedalus*, 136 Irving St., Cambridge, Mass. 02138 (617-491-2600).

The four remaining U.S. supercomputer centers, which are sponsored in part by the National Science Foundation, all issue newsletters and annual achievement summaries. The four are: the Cornell Theory Center (607-255-5193); the National Center for Supercomputing Applications at the University of Illinois (217-244-0072); the San Diego Supercomputer Center (619-534-5000); and the Pittsburgh Supercomputing Center (412-268-4960).

**PERFORMANCE.** A number of papers evaluating the performance of massively parallel processors and comparing the results with those of more conventional supercomputers have been published recently. "The performance realities of massively parallel processors: a case study" was written by O. M. Lubeck, M. L. Simmons, and H. J. Wasserman. It is available through e-mail address {mls,hjw,oml}@lanl.gov, or by writing to Los Alamos National Laboratory, Mail Stop B265, Los Alamos, N.M. 87545.

Two technical reports by the National Aeronautics and Space Administration's Ames Research Center also make interesting reading. "NAS Parallel Benchmark Results" is code number RNR-92-002; "Towards the Teraflops Capability for Computational Fluid Dynamics" is RNR-92-016. Both are available through the NAS Library, NASA Ames Research Center, Moffett Field, Calif. 94035; 415-604-4624.

"Supercomputer performance analysis and the Perfect Benchmarks" was published by the Association for Computing Machinery, New York City, in the proceedings of the 1990 International Conference on Supercomputing. It is also available as Technical Report No. 965 from the Center for Supercomputing Research and Development, University of Illinois, Urbana, Ill. 61801; 217-244-1608.

**ARCHITECTURES.** The report of the Purdue Workshop on "Grand Challenges in Computer Architecture for the Support of High Performance Computing" was issued in July. The report, No. 92-26, may be obtained by writing to Technical Reports, School of Electrical Engineering, 1285 Electrical Engineering Building, Purdue University, West Lafayette, Ind. 47907-1285; fax, 317-494-6440.

"Mathematical Foundations of High-Performance Computing and Communications" was released in October 1991 by the Board on Mathematical Sciences of the National Research Council, 2101 Constitution Ave., N.W., Washington, D.C. 20418; 202-334-2421.

**NATIONAL EFFORTS.** *High Performance Computing: Research and Practice in Japan*, edited by R. H. Mendez, was published earlier this year by John Wiley & Sons in New York City. A full issue of the *Japan Computer Quarterly* was devoted to "Real World Computing & Related Technologies" (No. 89, published earlier this year by Jipdec in Tokyo). Ulrich Wattenberg was the author of *Massively Parallel Optical and Neural Computing in Japan*, published by the IOS Press, Amsterdam, the Netherlands, 1992.

A transcript of testimony by Arvid G. Larson, the IEEE's Vice President-Professional Activities, on the competitiveness of the U.S. supercomputer industry is available through the IEEE Washington Office, 1828 L St., N.W., Suite 1202, Washington, D.C. 20036-5104; 202-785-8017. The testimony was given before the Legislation and National Security Subcommittee of the House Committee on Government Operations on July 1.

"Grand Challenges 1993: High Performance Computing and Communications" gives a broad view of the U.S. government's effort in this area, its goals, and the roles of the various agencies involved. The 68-page booklet, prepared by the President's Office of Science and Technology Policy, is available from the Federal Coordinating Council for Science, Engineering, and Technology,

Committee on Physical, Mathematical, and Engineering Sciences, c/o National Science Foundation, Computer and Information Science and Engineering Directorate, 1800 G St., N.W., Washington, D.C. 20550.

**APPLICATIONS.** For an in-depth analysis of what mathematics has to do with reality, the new book *The World Within the World*, by John D. Barrow (Oxford University Press, 1990), can be very enlightening.

"Aerodynamic Predictions for Supersonic Transport Aircraft" is but one example of the 150 or so recent summaries on the subject of computational fluid dynamics listed in the annual *Numerical Aerodynamic Simulation Technical Summary*. To get on the mailing list of the Numerical Aerodynamic Simulation Program, write to the NASA Ames Research Center, Moffett Field, Calif. 94035-1000.

There is an endless number of general articles on turbulence and fluid dynamics. Two of note are: "Morphological Structures Produced by Mixing in Chaotic Flows," by J. M. Ottino et al. (*Nature*, Vol. 333, no. 6172, June 2, 1988) and "Euler and Navier Stokes Computations for Two-dimensional Geometries Using Unstructured Meshes," by D. J. Mavriplis (ICASE Report No. 90-3, NASA CR-181977, January 1990).

**FUTURE TECHNOLOGIES.** For background information on the role of organic molecules in computation, see "The lure of molecular computing," by Michael Conrad, *IEEE Spectrum*, October 1986, pp. 55-60.

Details of how neural networks function are given by Robert Hecht-Nielsen in "Neurocomputing: picking the human brain," *Spectrum*, March 1988, pp. 36-41.

The novel physical effects encountered in the nanometer realm are described by Karl Hess and Gerald J. Iafrate in "Approaching the quantum limit," *Spectrum*, July 1992, pp. 44-49.

## Acknowledgments

*IEEE Spectrum* called upon many experts in preparing this special issue but is especially indebted to the following consultants for their advice (their identification with the report, though, should not be construed as an endorsement of it in its entirety): Malvin H. Kalos, Cornell Theory Center; Kenneth W. Kennedy (SM), Rice University; David J. Kuck (F), Center for Supercomputing Research and Development, University of Illinois at Urbana-Champaign; Sidney Karin (M) and Wayne Pfeiffer, San Diego Supercomputer Center; John P. Riganati (SM), Supercomputing Research Center; and Bruce D. Shriver (F), Consultant, Ossining, N.Y.

The issue editor would also like to acknowledge the assistance of Gary Smaby, Smaby Group Inc., and Jack Dongarra (M), University of Tennessee and Oak Ridge National Laboratory.



**For High Temperature  
Superconductor Answers,  
Go to the Source.**




# Are the new superconductors really practical for electronic devices?

**"YES, and they can give a significant performance advantage to your systems. At SUPERCONDUCTOR TECHNOLOGIES INC. (STI), the new high temperature superconductors (HTS) have developed rapidly, especially for microwave applications. Our high quality superconducting thin films, which have losses 100 times lower than copper or gold, can be patterned into thin-film circuits with Qs as high as 30,000. In fact, STI now provides a range of products that give extraordinary performance advantages to a number of industries.**

**"STI is now working with the microwave industry in an aggressive program to develop superior passive components with HTS. This includes resonators, filters, oscillators, delay lines and subsystems. In addition to HTS films, we provide custom HTS circuits, components and assemblies. And our HTS circuit fabrication department builds customers' designs from mask making to packaging.**

**"HTS theory is now HTS reality. Managers and systems engineers in a growing number of industries now have the opportunity to use the performance advantages of high temperature superconductors."**

*Jim Bybokas  
Vice President  
Superconductor Technologies*



*Examples (shown approximately 2X actual size) of custom HTS circuits include microstrip resonators, coplanar delay lines, lumped element inductors and microstrip filters.*

STI has been awarded more than 20 development contracts and has recently completed a major three-year program for DARPA to develop HTS materials and devices for microwave applications.

At STI, we have successfully developed HTS materials with transition temperatures over 100°K. As we advance the state of superconductor technology, we're continually adding innovative new products to the marketplace.

Make high performance HTS an integral part of your component and system development now by making



STI part of your development team. For more information, call Jim Bybokas, Vice President of Product Development at Superconductor Technologies Inc., at (805) 683-7646. Or write him at 460-F Ward Drive, Santa Barbara, CA 93111-2310... FAX (805) 683-8527.

You'll be surprised at what has been accomplished with high temperature superconductors!



**SUPERCONDUCTOR  
TECHNOLOGIES**

*Your complete HTS source.*



## Books

(Continued from p. 16)

tinct impression that technical and economic considerations determined, for the most part, the complex history of electricity in Chicago—a story revolving around utility magnate Samuel Insull and the rise of his Commonwealth Edison Co.

This story has been told well by Hughes in his award-winning *Networks of Power: Electrification in Western Society, 1880–1930* (Johns Hopkins University Press, Baltimore, Md., 1983), which has a chapter on Chicago, and by Forrest McDonald's 1962 biography of Insull. In many respects, *Electric City* is an amalgam of these two books plus a more recent one, David Nye's *Electrifying America: Social Meanings of a New Technology, 1880–1940* (MIT Press, Cambridge, Mass., 1990).

Platt, however, goes beyond Hughes and McDonald in several ways. He considers the history of coal, gaslight, and arc lighting in Chicago; provides a more in-depth account of municipal politics in the city; describes in more detail the growth of the Insull system from the center of the city to encompass a regional system of suburbs, small towns, and farms; and discusses the important topic of electricity in the home, which Hughes and McDonald neglect. Platt's analysis of the social implications of electricity,

however, does not match that of the path-breaking *Electrifying America*, which appeared too late to influence his work.

There are other shortcomings. Although Platt digs deeper into the sources of Insull's rate-making innovations and explains more fully his "gospel of consumption" (based on large-scale production of electricity, low rates for consumers, and high utility company profits), Hughes still gives a clearer account of Insull's influential system of electric utility economics.

On electricity in the home, Platt cites the respected thesis of historian Ruth Schwartz Cowan, to the effect that household technology eliminated the drudgery, but created "more work for mother" by replacing the work of men and boys and raising standards of living. But he dismisses the critical aspect of her thesis in support of his broader goal of overturning modern critiques of the United States' culture of consumption. The reader might also wonder why, if there was such a large "pent-up demand" for electricity in Chicago, Insull and Commonwealth Edison resorted to so many promotional schemes and massive advertising campaigns to sell electricity over the years. And, finally, Platt has missed an opportunity to compare, at least briefly, his wealth of information on Chicago with what we know about the electrification of such cities as New York, Bos-

ton, Detroit, Philadelphia, and Kansas City.

Despite these shortcomings, Platt has written an excellent book about an important and fascinating subject. His insights into the competition between gas and electricity, Insull's business and political accomplishments, and the cultural views of electricity place the book in front ranks of recent scholarship on the social meanings of technology and on those who invented and nurtured it.

Ronald R. Kline is an assistant professor of history of technology at Cornell University, Ithaca, N.Y., president of the IEEE Society on Social Implications of Technology, and author of *Steinmetz: Engineer and Socialist* (Johns Hopkins University Press, Baltimore, Md., 1992).



## Recent books

**Lenk's Laser Handbook: Featuring CD, CDV, and CD-ROM Technology.** Lenk, John D., McGraw-Hill, New York, 1992, 279 pp., \$39.95.

**Command-Level CICS Programming.** Varsegi, Alex, TAB/McGraw-Hill, Blue Ridge Summit, Pa., 1992, 282 pp., \$39.95.

COORDINATOR: Glenn Zorpette

Dr. Smith  
found  
two peaks...



Using PeakFit,  
Dr. Jones found  
three peaks...

## Guess who got the grant?

Because PeakFit™ uses sophisticated non-linear curve fitting techniques, it is far more effective than other data analysis tools in reducing noise and separating unresolved peaks! In fact, PeakFit is one of the most powerful software packages available today for chromatography and spectroscopy data analysis. For chromatographers, PeakFit extracts important system suitability information such as peak areas, asymmetries, and resolutions – plus column efficiency (number of theoretical plates) from raw data. For spectroscopists, PeakFit calculates peak amplitudes, areas, centers and widths – with complete uncertainty analysis, too!

Use PeakFit's 30 built-in peak functions (including Gaussian, Exponentially-Modified Gaussian, Voigt and many others specifically suited for chromatography and spectroscopy) or input and use your own! Output both graphs and data to over 150 printers, including PostScript®. Of course we can't promise that PeakFit will win you a grant (your competition may already be using PeakFit!), but we can promise that PeakFit will enhance your current analysis methods. Order PeakFit today and see just how many peaks there really are in your data!

Manufacturers: Call for OEM info!

New Version 3.0!


**PeakFit™**  
Peak Analysis Software  
for the IBM PC

For free brochure:  
800-874-1888 (U.S.)  
415-453-6700 FAX 415-453-7769

From the makers of SigmaPlot®!

**Jandel**  
SCIENTIFIC

2591 Kerner Blvd  
San Rafael, CA 94901  
Europe: Schimmelbuschstr. 25  
4006 Erkrath, Germany






# JOIN THE GLOBAL INTUITION NETWORK

The Global Intuition Network is supported by a grant from The Charlson Research Foundation. The purpose of The Network is to promote the applied use of intuition in decision making; share new knowledge on how to use this brain skill as it becomes known; and to promote ongoing research on intuitive processes for practical use in organizations of all types.

The focus of the Network is global, interdisciplinary, practical as well as academic, and on all levels and types of education. If you would like to join and participate in this Network (which is free), contact: Dr. Weston H. Agor, Director, Global Intuition Network, P.O. Box 614, UTEP, El Paso, Tx. 79968 or FAX 915-747-5111.





# Program notes

## Packing them in

Modern software cries out for more disk space. A scanned image of an 8 1/2-by-11-inch page requires at least a megabyte; a typical application now uses up at least 10 megabytes; and hefty applications like the new C/C++ compilers from Borland International Inc., Scotts Valley, Calif., and Microsoft Corp., Redmond, Wash., absorb 30-50 megabytes.

A 100-megabyte hard disk—standard equipment with most new 386- and 486-based personal computers—is just not big enough for most users. To make matters worse, prices for hard disks are currently going up, not down. So to pack more information into their disk space, users are starting to compress and decompress their files.

The software that does this uses one algorithm to compact files for storage and another to restore them. Bulletin boards and shareware distributors, for example, have long used data compression programs to pack (compress) files, so that they become quicker to upload and download. Users later



*This image was created by Image Incorporated compression software from Iterated Systems Inc., Norcross, Ga. and illustrates how little is lost by lossy algorithms. The right half of the picture is the original digitized image. The left half was compressed into an FIF file of 10 kilobytes or so and then decompressed.*


unpack the files on their own PCs, using a data decompression program.

The algorithms fall into two categories: lossless and lossy. The lossless type must be used if the restored file has to look exactly like the original. For instance, algorithms for compressing and decompressing programs and text files must be lossless. Lossy algorithms suffice when the size of a compressed file is more important than exact restoration.

Programs for storing images and sound use one of several lossy algorithms, trading reduced size for imperfect image reproduction. Still images may be stored in the Joint Photographic Experts Group (JPEG) format or in the generic compact disk-interactive (CD-I) format; images and audio may be stored in Intel's proprietary DV-I (digital video-interactive) format, and motion images may be stored in Motion Picture Experts Group (MPEG) format. All the algorithms in these techniques compress files to about a fortieth of their original size with only minor losses after decompression.

*(Continued on p. 83)*

## Make the right move!



let **MRJ** guide you  
through the  
supercomputing maze.

**MRJ** is an advanced technology and engineering services company with unique expertise in supercomputing. Our satisfied customers include Fortune 500 companies and Government agencies.

**MRJ** has a broad range of parallel and distributed hardware and software expertise. Our independence allows us to objectively determine the most cost-effective platform and software solution to satisfy your specific needs.

In 1986, **MRJ** became the first commercial buyer of the Connection Machine™. Since then, we have specialized in the development of operational supercomputing applications and software on many platforms. **MRJ** has developed applications in:

- Information Search and Retrieval
- Link Analysis
- Intelligent Decision Systems
- Simulation, Modeling, and Optimization
- Signal and Image Analysis
- Electromagnetic Scattering

Contact: Mr. Andre Lachance - MRJ, Inc.

10455 White Granite Drive, Oakton, VA 22124

(703) 218-0329

super@mrj.com

Fax (703) 385-4637

MRJ is always looking for creative, motivated professionals seeking a challenge. Send resume - ATTN: Personnel-92-106

U.S. Citizenship Required - Current SSBI Preferred

Connection Machine is a trademark of Thinking Machines, Inc. All other brand and product names are trademarks of their respective companies.



## READER GUIDE TO PRODUCTS AND SERVICES

### VOLUNTEER OPPORTUNITIES IN PRECOLLEGE EDUCATION

IEEE United States Activities has updated and reissued its "Directory of Volunteer Opportunities in Precollege Mathematics and Science Education for Engineers and Scientists." The new edition was printed jointly by IEEE-USA and the Engineers for Education Program.

The directory lists 34 extracurricular programs that seek to improve and enhance math, science, and technology education for elementary and secondary school students. The programs are nationally coordinated but operate locally in various cities. All are heavily dependent upon volunteers for their human resources.

Complimentary copies of the new directory are available on request from the IEEE-USA Office. Call 202-785-0017 or CIRCLE #81 on the Reader Service Card.

### A Comprehensive 3-Day IEEE Standards Seminar

New Orleans, La. November 2-4, 1992

#### Industrial and Commercial Power Systems Newly Revised and Expanded!

Developed specifically for today's electrical and electronics engineers, this seminar focuses on the safe and effective design, operation, and protective maintenance of industrial and commercial power systems.

#### Course Overview

In this seminar, you'll learn about:

- Industrial power systems design, configurations, grounding, and voltage considerations
- Commercial building design, power distribution, voltage considerations, distribution systems, and high-rise electrical systems
- Equipment selection and application
- Equipment protection (transformer, bus, cable circuit, motor, and generator)
- Power systems protection and principles of coordination
- Savings with energy conservation management, electrical rates and costs, co-generation, and metering and load management
- Systems analyses and specialty studies
- Short-circuit calculations (formal vs. fast manual)
- Emergency power types, applications, and ratings; on-site generator grounding
- Electrical safety testing and preventive maintenance
- Power systems reliability/availability considerations

Fee: \$990 IEEE Members: \$940

The IEEE registration fee includes continental breakfast, lunch, and refreshments.

For more information, CIRCLE #82 on the Reader Service Card.

### New IEEE Videoconference Program October 29, 1992

HOW WILL WE BROADCAST HDTV?  
Recent advances in the technology of data compression

reduction have changed our ideas about the transport of television. This videoconference will review how some of the most important compression algorithms can be used in your applications.

We can now imagine carrying full-motion video on ordinary telephone wire pairs or transmitting digital high-definition television (HDTV) either over the air or as an integrated multimedia service on fiber. All of these services and systems rely upon a common base of digital video signal-processing techniques.

This program is critical viewing for the working engineer and should not be missed by managers and technicians in this field.

Lead Presenter: Jules A. Bellisio, executive director, Video Systems and Signal Processing Department, Bellcore, Red Bank, N.J.

Corporate rate (single site): \$1800 per broadcast.

University rate (single site): \$850 per broadcast.

Questions: Call Judy Brady, IEEE Marketing, 908-562-3991.

For more information, CIRCLE #83 on the Reader Service Card.

### STUDENT EMPLOYMENT GUIDE

IEEE United States Activities has just published the second edition of the student version of *The Employment Guide for Engineers and Scientists*. In this two-volume publication, the first volume includes basic information on conducting a job search, writ-

(Continued overleaf)

## READER SERVICE (CIRCLE NUMBERS)

### 1 PRODUCT INFORMATION

1	10	19	28	37	46	55	64	73	82	91	100	109	118	127	136	145	154	163	172
2	11	20	29	38	47	56	65	74	83	92	101	110	119	128	137	146	155	164	173
3	12	21	30	39	48	57	66	75	84	93	102	111	120	129	138	147	156	165	174
4	13	22	31	40	49	58	67	76	85	94	103	112	121	130	139	148	157	166	175
5	14	23	32	41	50	59	68	77	86	95	104	113	122	131	140	149	158	167	176
6	15	24	33	42	51	60	69	78	87	96	105	114	123	132	141	150	159	168	177
7	16	25	34	43	52	61	70	79	88	97	106	115	124	133	142	151	160	169	178
8	17	26	35	44	53	62	71	80	89	98	107	116	125	134	143	152	161	170	179
9	18	27	36	45	54	63	72	81	90	99	108	117	126	135	144	153	162	171	180

### MEMBERSHIP INFORMATION

300 Regular 301 Student

Print or Type only

Name \_\_\_\_\_ Title \_\_\_\_\_

Company \_\_\_\_\_

Address \_\_\_\_\_

City \_\_\_\_\_ State \_\_\_\_\_ Zip \_\_\_\_\_

Country \_\_\_\_\_ Business Phone \_\_\_\_\_

If this represents a change in address, please fill out form on p. 9

### 3 ADDITIONAL COMMENTS

I would like: \_\_\_\_\_

DELIVERY: I received this issue \_\_\_\_\_ (date).

For faster information, fax this card to 413-637-4343

Void after December 1, 1992

Void overseas January 1, 1993

## READER FEEDBACK (CIRCLE NUMBERS)

### 2 EDITORIAL MATTER

Article/Item	Comment		
	No Satisfied	Just OK	Too Satisfied
Supercomputers			
Massive parallel processing	01	02	03
Software	04	05	06
Visualization	07	08	09
Networks/storage	10	11	12
Japan	13	14	15
Applications	16	17	18
Devices	19	20	21
Packaging	22	23	24
Future	25	26	27
Newslog	28	29	30
EEs' tools & toys	31	32	33
Program notes	34	35	36

Did you find this material stimulating or thought-provoking?

Item	Yes	No
Spectral lines	37	38
Forum	39	40
Books	41	42
Reflections	43	44
Engineer at large	45	46



## READER GUIDE TO PRODUCTS AND SERVICES

ing a résumé, preparing for interviews, and developing a network to assist in the job search process.

The second volume is a *Directory of Employers of IEEE Members*, a state-by-state sourcebook that lists hundreds of companies, and includes telephone numbers and contact persons.

The book may be ordered through the IEEE Service Center (Order Number UH0188-3) for \$14.95 for members and \$19.95 for nonmembers, plus tax and shipping. Call 1-800-678-IEEE to order.

For more information, CIRCLE #44 on the Reader Service Card.

### NEW IEEE VIDEO

#### Wireless Communications

Sponsored by Virginia Polytechnic Institute and State University's Mobile & Portable Radio Research Group (MPRG) and the Division of Continuing Education, this new video was made in cooperation with the IEEE Communications Society and the IEEE Vehicular Technology Society. Presenters: Donald Cox of Bellcore, David Goodman of Rutgers University's Winlab, Larry Milstein of the University of California at San Diego, John Proakis of Northeastern University, Don Schilling of City College of New York, D.R. Vaman of Stevens Institute of Technology, Joseph McGeehan of Bristol's Center for Telecommunications Research, Allen Salmasi of Qualcomm, Bruce Tuch of NCR, Tien Hou of AT&T, Chandos Ripinski of LACE, Greg Vatt of Motorola, Roger Newell, publisher of *Microcell Report*, and Stuart Meyer, past president of the IEEE Vehicular Technology Society.

In this symposium, standards body participants,

researchers from industry, and engineering faculty from around the globe represent the broad range of interests and activities in wireless communications, including: radio local-area networks (RLANs), personal communications networks (PCNs), digital cellular radio, and low earth orbit (LEO) satellites. Recent technological advances in each of these areas, and how these different systems will compete with or complement each other in the future, will be some of the many topics considered during this video symposium.

Product No.	ISBN	Member	List
HV0239-4	0-7803-0318-0	\$295.00	\$395.00
Recording date	Run time	Standard	
6/91	7-12 hours/	NTSC, PAL	
	4-6 tapes		

To order, call 1-800-678-IEEE; fax, 1-800-981-9667. Outside the United States, call 908-981-0060.

For more information, CIRCLE #85 on the Reader Service Card.

#### APPLICATIONS OF MATHEMATICS IN ELECTRICAL ENGINEERING

This IEEE Educational Activities Department release is a highly applications-oriented, self-study review course in mathematics for the electrical and electronics engineer. Covering algebra, trigonometry, calculus, and complex arithmetic with an emphasis on linear equations, this course is intended for engineers with general mathematical backgrounds who wish to enhance their skills in direct applications and become more familiar with the associat-

ed computational methods.

A program disk containing many complex computations serves as a valuable reference tool providing long-term support. The course contains a study guide, answer book, program disk, final exam, and a textbook, *The Calculus Tutoring Book*, by Carol and Robert Ash, IEEE Press.

For more information on this course and other continuing education materials, call 908-562-5498, or CIRCLE #86 on the Reader Service Card.

#### COMPETING IN THE WORLD MARKETPLACE

IEEE United States Activities recently published *How the United States Can Compete in the World Marketplace*, a practical guide developed by the Committee on U.S. Competitiveness.

The 24-page booklet outlines several steps that the United States can take to regain its competitive edge, such as creating more funds for private capital investment; modifying Federal antitrust policies to conform with the current global market situation; investing in long-term manufacturing processes; employing managers with a good understanding of the technologies and economics of their businesses; incorporating engineering courses in product development, manufacturing technology, and systems engineering into school curricula; and instituting pre-competitive R&D geared toward supporting a product line rather than specific products. The booklet is available free of charge from the IEEE-USA Office in Washington, D.C.

To order, CIRCLE #87 on the Reader Service Card.



NO POSTAGE  
NECESSARY  
IF MAILED  
IN THE  
UNITED STATES

### BUSINESS REPLY MAIL

FIRST CLASS MAIL PERMIT NO. 885, PITTSFIELD, MA

POSTAGE WILL BE PAID BY ADDRESSEE

IEEE  
**SPECTRUM**

READER SERVICE MANAGEMENT DEPT.  
PO BOX 5149  
PITTSFIELD, MA 01203-9740





## Program notes

(Continued from p. 80)

A newer, computation-intensive compression algorithm reuses and upgrades existing images. Dubbed the fractal-image format (FIF), it converts images into a series of overlaid fractal shapes that can be stored in mathematical form. The images shrink to an eightieth of their former selves, again with only minor losses after decompression.

FIF is currently being used in a software package called Image Incorporated. Because decompression is independent of display resolution, images captured in an FIF file from a comparatively low-resolution medium can be decompressed to fill higher-resolution screens or other media. This means that VGA images may be converted to SuperVGA or XGA without the graininess that mars enlargement of files compressed by other formats.

Still, DV-I or CD-I images may be packed or unpacked in real time with a 33-MHz 80386 or a 25-MHz 68040, while FIF image compression takes much longer, hours rather than minutes. Application developers need the computing power of a 486 or a dedicated reduced-instruction-set processor to compress many images. However, decompressing an FIF file typically takes seconds, so applications using precompressed

FIF images are as fast as applications using DV-I or CD-I precompressed images. For more on Image Incorporated, contact: *Iterated Systems Inc., 5550-A Peachtree Parkway, Norcross, Ga. 30092; 404-840-0310; or circle 100.*

### From hype to hope

Multimedia—or rather, the growing enthusiasm for using it on personal computers—is behind the new popularity of lossy compression techniques. What multimedia needs now is a single standard. Microsoft Corp., Redmond, Wash., began this process last October when it introduced its specification for an IBM-compatible multimedia personal computer, the MPC. Now that Windows 3.1 is here [*IEEE Spectrum*, June 1992, p. 58], the MPC specification is slightly clearer. Text, images, and music are stored on a CD ROM, while a high-resolution graphics card is used for video and a digital audio card for sound. Of course, all the elements are linked by the Windows operating environment; the various compression algorithms are installed as drivers, much as for printers.

To learn more about Microsoft's vision for multimedia, get hold of its trilogy, *Windows Multimedia Reference Library*. Contact: *Microsoft Corp., 1 Microsoft Way, Redmond, Wash. 98052; 206-882-8080; or circle 101.*

### Kidnapped: no Treasure Island

The recent kidnapping of John Warnock, the creator of the pioneering desktop publishing software, PostScript, and the founder of Adobe Systems Inc., illustrates how much the software business has changed. Ten years ago, the founders of Adobe, Aldus, Microsoft, and hundreds of other software companies were struggling young technologists who lived anonymous lives. Today, they are rich, middle-aged celebrities.

Warnock's ordeal obeyed the standard Hollywood script for the abduction of an industrial tycoon: he was held for several days, then released once a ransom was paid, and the kidnappers were caught soon after. The only surprise in the affair was the smallness of the ransom. The kidnappers apparently felt that Warnock had lost so much money to Microsoft in the TrueType/Adobe Type Manager font war that a budget ransom of US \$650 000 was appropriate.

Warnock's tale and stories of other software millionaires are the purview of Robert X. Cringely's recent book, *Accidental Empires* (Addison-Wesley, 1992).

CONTRIBUTOR: John R. Hines is silicon sensors engineer at Honeywell Inc.'s MicroSwitch Division, Richardson, Texas.

COORDINATOR: Richard Comerford

## AT CONVEX SUPERCOMPUTING & SUPER OPPORTUNITIES START WITH CUSTOMER FOCUS.

Customer focus is the key to the **CONVEX Computer Corporation** environment, where some of the brightest minds are creating supercomputing solutions for world-class challenges. CONVEX has been redefining the opportunities in affordable supercomputing since 1982, when we introduced the first supercomputer under \$1 million. Now, we're building on our tradition of leadership with new generations of high performance vector/parallel systems and massively parallel processing systems. Opportunities currently exist in the following areas:

- Software Development
- Diagnostics Development
- S/W Tools Development
- I/O Engineering
- Networking Software
- Compiler Development
- Hardware Development
- Technical Writing
- Software Test
- Electronic Packaging
- Operating Systems
- Systems Performance Engineering

For immediate consideration, please send your resume to: **Human Resources, CONVEX Computer Corporation, Dept. IEEE, P.O. Box 833851, Richardson, TX 75083-3851.** We are an equal opportunity employer, M/F/D/V. Principals only, please.



**CONVEX**

## Analog Circuit Simulation

Integrated

Powerful

- Integrated Schematic Entry
- Extensive Device Libraries
- Powerful SPICE Simulator
- Monte Carlo Analysis
- Easy to use Waveform Processing
- Training Classes

Intusoft has it all at an Affordable Price!

Affordable

The ICAPS simulation system allows an engineer to enter a circuit into the computer and evaluate its behavior before actually building the circuit. It includes 4 integrated modules. **SPICENet** is a schematic entry program that generates a complete SPICE netlist and alleviates many of the headaches associated with older SPICE programs. **PRESPICE** adds extensive libraries with over 1200 parts, as well as the ability to add your own models. The **ISPACE** module runs on all PC computers and performs the actual AC, DC, time, noise, Fourier, and temperature analyses. Special extended RAM versions capable of simulating large circuits are available. The last module, **INTUSCOPE**, displays and measures the ISPACE output data. Starting at \$95 for ISPACE, complete systems are available for under \$1000.

For Information and Your Free Demonstration Kit  
Write Intusoft P.O. Box 710 San Pedro, CA 90733-0710  
or Call (310) 833-0710, FAX (310) 833-9658

**intusoft**


Circle No. 5



# EEs' tools & toys

## Graphics calculator finds polynomial roots

Texas Instruments Inc.'s latest scientific calculator is a programmable graphics unit that performs such feats as: finding all the roots of a polynomial of up to 30th order at the touch of a single button; solving systems of up to 30 simultaneous equations; handling matrices with up to 255 by 255 elements; and determining the value of any variable in a single equation.



Equation solving is the forte of the TI-85 calculator, which can find all the roots of a 30th-order polynomial at a stroke.

Of course, it also performs all of the usual scientific calculator functions—arithmetic, trigonometry, statistics, and so forth—and plots the results.

The calculator's graphing capability is useful for getting a sense of the overall behavior of a function and for rapidly estimating its roots. It also makes clear where plotted functions intersect. A "trace" feature allows the user to slide a cursor along a graph while a numeric display is giving the cursor's precise coordinates.

TI-85 programs may range from a sequence of simple commands to lengthy routines with branches, loops, and I/O instructions. The number and complexity of user programs is limited only by the calculator's 32K bytes of RAM.

Worth noting is the 85's ability to intermingle rectangular and polar forms of complex numbers. It can be told to add  $5.20 + 4.50j$  and  $6.00 \angle 0.50$  rad, for example, and immediately produce the result  $10.47 +$

$7.38j$ , or  $12.80 \angle 0.61$  rad, depending on which display mode has been set as the default.

The calculator has an I/O port that may be used to establish links with a personal computer or another TI-85. A 760-mm cable for interconnecting two TI-85s comes with the calculator. The software and special cable needed for connection to a personal computer will be available at an extra charge later this year.

The TI-85 measures 175 by 81 by 20 mm and weighs 0.240 kg. It comes complete with a hard slide-on cover, four AAA cells plus one lithium cell to save programs and data while the main battery is being changed, and a 352-page guide book. It carries a suggested retail price of US \$130. *Contact: Texas Instruments Inc., Consumer Relations, Box 53, Lubbock, Texas 79408-0053; 800-TI-CARES; or circle 102.*

## SOFTWARE

### Excel-lent design tools

Although not everybody realizes it, electronic spreadsheet programs are useful for solving a variety of engineering problems. With the appropriate worksheets and macros, they can calculate and plot filter responses, spectra, and transfer functions; find roots of equations; perform number base conversions; generate ROM logic; and so on.

According to the folks at Engineering Solutions, Microsoft Excel, "because of its exceptional power in matrix, logical, mathematical and graphical functions," is a spreadsheet peculiarly well suited to this type of work. So they have put together a collection of worksheets and macros based on Excel and are offering them for sale individually at prices ranging from \$15 to \$50. With few exceptions, the routines may be had in both PC and Macintosh formats.

Currently available software is described in the Winter 1992 catalog, which is offered free of charge. The company is also interested in acquiring new routines, for which it will pay on a royalty basis. *Contact: Engineering Solutions, Box 570159, Tarzana, Calif. 91356; 818-772-7231 or circle 103.*

## GENERAL INTEREST

### Tetris set to music

Some things are harder to understand than others. It makes sense that Spectrum HoloByte has come out with a new version of the popular computer game Tetris. After all, what good is a game without VGA graphics

and the ability to work with the latest sound boards? And what kind of nerd would play Tetris without accompanying selections from Glinka's opera *Russian and Ludmilla*?

Equally easy to comprehend is the necessity of supporting Novell Netware (or a simple null modem cable) for head-to-head competitive or cooperative play.

What strains the intellect and keeps one up nights is something else. Why did the marketing geniuses behind the product choose to name this new and improved program "Tetris Classic"? Wouldn't "Tetris 92" or "Tetris Moderne" or even "Turbo Tetris" have been more appropriate? It's hard to say. The first reader to write in with a plausible explanation will receive two used tickets to a Mostly Mozart concert in New York City at which no Russian music at all was played.

Oh yes, the program is available through the usual retail outlets at a suggested price of \$49.95. And a high-resolution Windows version, priced at \$44.95, should be out by the time this hits print. *Contact: Spectrum HoloByte, 2061 Challenger Dr., Alameda, Calif. 94501; 510-522-3584; fax, 510-522-3587; technical information, 510-522-1164; to place an order, 800-695-GAME; or circle 104.*

## AUTOMATION

### Tiny, low-cost PLC

Aimed at small control jobs, the Little PLC is a very compact (110 by 72.4 mm) programmable logic controller with a price tag of only \$195. It has eight optically isolated inputs, eight relay-driver outputs, its own switching power supply, and an RS-485 serial communications port. Other standard features include a watchdog timer, power failure detection, a battery-backed time-and-date clock, and battery-backed RAM.

Instead of ladder logic, the Little PLC is programmed using Z-World's Dynamic C software—a package consisting of a compiler, an editor, and a debugger—which runs on a PC. After compilation, programs developed on the PC are downloaded to the Little PLC over its serial port.

Like the hardware, the software package is priced at \$195, making the total system price just \$390. Expansion boards are also available. *Contact: Z-World Engineering, 1724 Picasso Ave., Davis, Calif. 95616; 916-757-3737; fax, 916-753-5141; or circle 105.*

*COORDINATOR: Michael Riezenman  
CONSULTANT: Paul A.T. Wolfgang, Boeing Defense & Space Group*



## Forum

(Continued from p. 22)

*tion Handbook* is a major contributor to misinterpretations and misapplications of this standard. This problem has been recognized and is being addressed by the Defense Software Development Standards Advisory Board (DSSAB). While document-driven approaches and generation of paperwork is still common on some projects, the DOD-STD-2167A specifically encourages transition to automated software development libraries and views documentation as a routine by-product of the engineering process.

As a matter of fact, the DOD-STD-2167A has been a major driver behind the rapid progress in the CASE industry. The standard establishes a sizeable market segment where a uniform set of standards and practices are contractually enforced. This facilitates the development of CASE tools and Software Engineering Environments (SEEs). While voluntary standards have contributed much to software engineering, they are relatively silent in the area of specifics defining the data elements necessary to develop, test, and manage the software. It is in this critical area of software engineering that the DOD-STD-2167A and its integrated SDL-based Data Item Descriptions (DIDs) have had their most significant impact on the CASE industry. Practically all major CASE vendors claim compliance with the DOD-STD-2167A process and provide templates for standard document types.

Contrary to the impression left by the article, the defense software community is a major factor in the development of international voluntary standards. It was instrumental in forming the USA Technical Advisory Group for what is now ISO/IEC Subcommittee 7 in 1984. It provided the subcommittee's first chair and is actively engaged in its leadership. Significant progress has been made since 1989 to harmonize defense and voluntary (that is, IEEE) standards so as to establish a unified U.S. position in the international marketplace.

Ole Golubjatnikov  
Syracuse, N.Y.

### Where age counts

I was assured by an American professor of engineering that a Microsoft advertisement recently published in the *London Times* would have been illegal in the United States of America.

Why would it be illegal? Because it is a clear example of age prejudice, with its three age ranges of 27-38, 27-35, and 24-34.

How rapidly the human mind decays! *Sic transit gloria mundi*, at least in Britain, but not, thank goodness, in the United States!

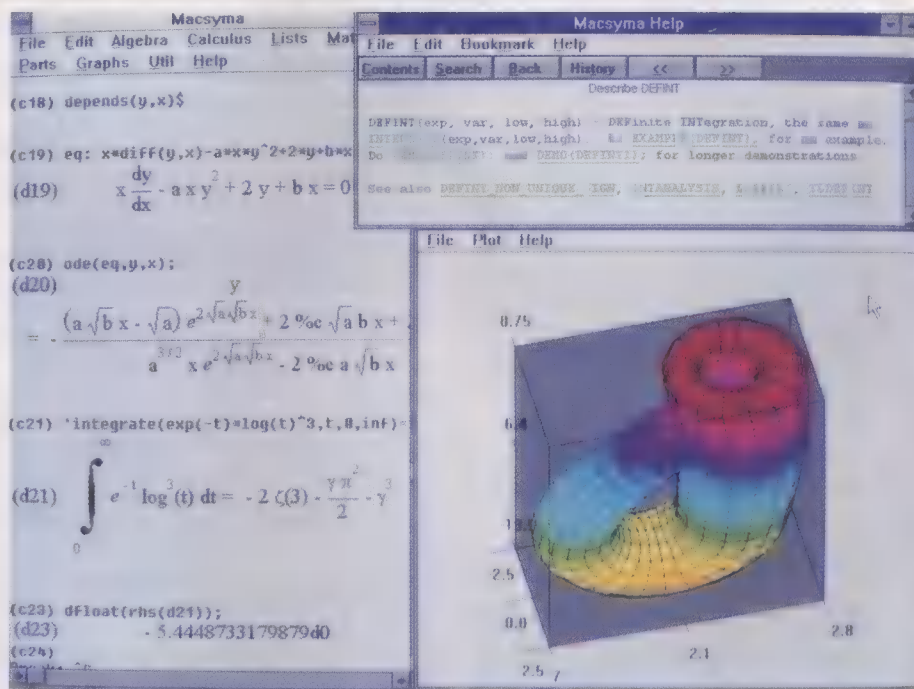
I find the advertisement to be as nasty and

unpleasant a piece of age discrimination as you could hope to find. It tells us what kind of company Microsoft really is, which is to say an unpleasant collection of age-prejudiced bigots.

Although I write on my own behalf, as a member of the Council of the British Institution of Electrical Engineers, I am in contact with some of the membership of the British engineering profession and believe that I am not alone in finding Microsoft contemptible.

A. Sandman  
London, United Kingdom

Readers are invited to comment in this department on material previously published in *IEEE Spectrum*; on the policies and operations of the IEEE; and on technical, economic, or social matters of interest to the electrical and electronics engineering profession. Short, concise letters are preferred. The Editor reserves the right to limit debate on controversial issues. Contacts: Forum, *IEEE Spectrum*, 345 E. 47th St., New York, N.Y. 10017, U.S.A.; fax, 212-705-7453. The Comppail address is ieeeespectrum. The computer bulletin board number is 212-705-7308 and the password is SPECTRUM; for more information, call 212-705-7305 and ask for the Author's Guide.



## The Superstar of Mathematical Software

**Macsyma** Now Costs \$349.00\*

### An Unbeatable Value on PCs and Workstations

Macsyma not only offers the most complete and reliable portfolio of mathematical capabilities, it now has an unbeatable price. Macsyma's new PC version runs as a regular Windows 3.1 application with many improvements.

- World's best trig simplification and best differential equation packages are now even stronger. Factoring of large integers is 1,000 times faster.
- Color Scientific Graphics includes graphic windows with user controls for rotate, zoom, hardcopy printing, postscript output, colors, plot labels. Also, plotting of parameter surfaces, improved hidden line removal.
- New screen display of special math symbols, exponents, and Greek letters makes output more readable.†
- Easy-to-use hypertext† on-line help system with 1,500 text entries, 500 executable examples and demonstrations. Function templates† indicate function names and argument slots.

If you want the best mathematical software at a price the competition can't touch, call 1-800-MACSYMA (1-800-622-7962) or FAX 617-646-3161. Quantity discounts and academic prices are available. \*Small workstations start at under \$1,000.


†only on PC at this time.

## Macsyma Inc.

20 Academy Street, Arlington, MA 02174

Circle No. 26





"All of the C language routines you need to write an impressive scientific graphing program of your own. Highly Recommended."

—PC Magazine

#### GraphiC Features

- Native Tektronix<sup>®</sup> format
- 3-D curves & surfaces
- Shaded contour plots
- Linear, log, polar, error function plots
- Smith, bar, pie charts
- Color PostScript<sup>®</sup> & EPS
- Export HPGL, HPGL/2, TIFF, PIC, GEM, Tek 4105
- TIGA, DGIS, 8514 and super EGA/VGA support
- Source code included
- 22 fonts including Greek, math, and European characters
- 9 line types, 16 symbol types, error bars, . . .
- Macintosh<sup>®</sup> and Phar Lap<sup>®</sup> 286/DOS-Extender versions available

Licensed for personal use only

### Scientific Endeavors

508 North Kentucky St., Kingston, TN 37663 USA  
(615) 376-4146 FAX: (615) 376-1571

Circle No. 9

## Reading worth writing for.

If you're looking for some good reading, you've just found it. The free Consumer Information Catalog.

The Catalog lists about 200 federal publications, many of them free. They can help you eat right, manage your money, stay healthy, plan your child's education, learn about federal benefits and more.

So sharpen your pencil. Write for the free Consumer Information Catalog. And get reading worth writing for.



Consumer Information Center  
Department RW  
Pueblo, Colorado 81009

U.S. General Services Administration.

## Engineer at large

(Continued from p. 18)

Electric Co. and as a member of the governing council of the academy. The award recognizes "statesmanship in the field of technology," as well as active involvement in determining science and technology policy, promoting technological development, and contributing to industry-Government-university relationships.

The awards each consist of a gold medalion and a certificate. They will be presented at the academy's annual meeting to be held Sept. 29.

### Public policies on three issues

At the American Association of Engineering Societies (AAES) in Washington, D.C., the new Engineers' Public Policy Council appointed three task forces in July to explore the development of policy positions in three areas: competitiveness, the environment, and R&D. Each task force has a lead society to provide staff support: the IEEE for competitiveness, the American Institute of Chemical Engineers for the environment, and the American Society of Mechanical Engineers for R&D.

Additionally, the council approved its organizational charter, which was subsequently ratified by the AAES board of governors in early August; began planning for its 1993 Government Affairs Conference, to be held March 3; and initiated efforts to identify leadership positions in the Federal government that should be filled by individuals with technical backgrounds.

According to its charter, the Engineers' Public Policy Council is to "serve as the principal forum within AAES for the coordination and development of . . . policies affecting both the professional and technical practice of engineering" and to provide "technical information and engineering judgment on a wide range of issues to public policymakers, coalitions, appropriate organizations, and the public at large."

### Symposium on smaller defense budget

Those concerned with the downsizing of the defense budget by the U.S. government may care to attend a symposium this month organized by the IEEE-USA Technology Policy Council, "Coping with the impact of defense budget restructuring." It takes place Sept. 15 and 16 at the Stouffer Concourse Hotel in Arlington, Va. Speakers will be from the departments of Defense and Commerce, the Congressional Budget Office, industry, academia, and the IEEE. Contact: Chris Brantley, IEEE-USA, 1828 L St., N.W., Suite 1202, Washington, D.C. 20036; 202-785-0017; fax, 202-785-0835.

COORDINATOR: Alfred Rosenblatt

## ASSOCIATE GENERAL MANAGER

### VOLUNTEER SERVICES

The Institute of Electrical and Electronics Engineers, Inc., the world's largest technical professional society, is seeking candidates for this executive staff position which is headquartered in Piscataway, NJ.

Reporting to the General Manager, the Associate General Manager for Volunteer Services will manage the departments which support the membership and volunteer structure, including Technical, Professional, Regional and Educational Activities.

The position is responsible for providing high quality services to the volunteer organization and maintaining close working relationships with the volunteer leadership, including the IEEE technical societies.

The position calls for an executive level manager with an appropriate engineering degree who combines strong technical and management background with outstanding human relations skills, and who has demonstrated an ability to work successfully with the senior level of volunteer organizations.

For consideration, please send your confidential resume, including salary history to:

Human Resources Director,  
The Institute of Electrical and Electronics Engineers, Inc., 345 East 47 St., New York, NY 10017.



IEEE

AN EQUAL OPPORTUNITY EMPLOYER M/F/H



# PRODUCT PROFILES

*The **PRODUCT PROFILES** advertising section provides a distinctive focus on both new and commercially established products and services that offer practical solutions for engineering and scientific professionals in today's fast-paced, rapidly changing high-tech environment.*

For rates and production information, phone 212-705-7579

## JAPAN NEW MATERIALS REPORT

Japan has rapidly emerged as the world leader in the development, production and application of an expanding range of advanced materials, particularly those for electronics applications.

**Japan New Materials Report** provides full coverage of the technological and business developments in this dynamic industry. Materials covered include single crystals, fine ceramics, ultra-pure metals, metal alloys, engineering plastics, adhesives, composites, functionally gradient materials and superconductors.

Published bimonthly. U.S.\$83 per year (\$150 for two years or \$210 for three years). Price includes airmail delivery.

Please make check payable to: **Japan New Materials Report**, P.O. Box 4129, Bellevue, WA 98009-4129, USA.

Name

Title

Company

Address



**\$99** each

**SIGLAB Signal Lab**  
Digital Filter Design  
DSP  $\mu$ P Code Generators

**DSP Software**  
(800) 741-7440


The Athena Group, Inc.  
3424 N.W. 31st Street, Gainesville, FL 32605 USA  
Tel: (904) 371-2567 FAX: (904) 373-5182  
Export pricing higher

Circle No. 50

## HERMETICALLY SEALED

**Mini Connectors, Headers and Terminals**


**Pressure and Vacuum Types**



**DETORONICS**

10660 E. Rush St. • El Monte, CA 91733-3432  
Phone (714) 579-7130 • FAX (818) 579-1936  
Send your drawing or give us a call.

Circle No. 51



## Polysoft TechPlot



The first powerful technical plotting and data processing package completely in graphics mode for DOS. **TechPlot** features a new graphical interface that is faster and more intuitive. **TechPlot** offers many plot types and has an on-screen plot editor that makes it easy to design and modify graphs. **TechPlot** also provides a comprehensive data sheet that performs complete statistical analysis, math transformation, digital data processing, and nonlinear curve fitting. List: \$449. Introductory: \$299.

**Polysoft**

P. O. Box 526368, Salt Lake City, UT 84152, USA  
Tel: (801) 485-0466 FAX: (801) 485-0480

Circle No. 53

## Analog & Digital Circuit Design & Simulation for the Macintosh & PC with Windows



**Beige Bag Software** Ph: (313) 663-4309  
Fax: (313) 663-0725 • 715 Barclay Ct. Ann Arbor, MI 48105

Circle No. 54

## Instant Microcontroller



## Instant C Programming

Don't use a microprocessor, use a **SmartBlock™** microcontroller module to build your custom controller. Our \$195 interactive **Dynamic C™** development system makes programming easy. 3.5 x 2.5 inch module includes microprocessor, memory, time/date clock, eeprom, watchdog, serial ports and more. As low as \$59 in quantity. The efficiency of a custom design without the headaches.

**Z-World Engineering**

1724 Picasso Ave., Davis, CA 95616 USA

Tel: (916) 757-3737

Regular Fax: (916) 753-5141

Automatic Fax: (916)-753-0618

(Call from your fax, request data sheet #14.)

Circle No. 52



# Employment opportunities

Organizations seeking engineers and scientists describe their various openings in the following advertising section. In order to conform to the Age Discrimination in Employment Act and to discourage age discrimination, IEEE may reject any advertisement containing any of these phrases or similar ones, "recent college grads," "1-4 years maximum experience," "up to 5 years experience," or "10 years maximum experience." IEEE reserves the right to append to any advertisement, without specific notice to the advertiser, "Experience ranges ■■■ suggested minimum requirements, not maximums." IEEE assumes that, since advertisers have been notified of this policy in advance, they agree that any experience requirements, whether

stated as ranges or otherwise, will be construed by the reader as minimum requirements only. While IEEE does not ban the use of the term "entry level," its use is discouraged since, to some, it connotes an age rather than an experience designation. IEEE accepts employment advertising to apprise its members of opportunities. Interested parties should be aware that the political and humanistic values of certain advertisers may differ from their own. IEEE encourages employers to offer salaries that are competitive, but occasionally a salary may be offered that is significantly below currently acceptable levels. In such cases the reader may wish to inquire of the employer whether extenuating circumstances apply.

To place an advertisement in Spectrum's Employment Opportunities section, contact the nearest Spectrum sales office

New York	Boston	Chicago	San Francisco	Los Angeles	Atlanta	Dallas
212-705-7760	508-255-4014	708-446-1444	415-386-5202	310-649-3800	404-256-3800	214-553-9896

For production/mechanical information contact Theresa Fitzpatrick Advertising Production Manager, 212-705-7579

IEEE Spectrum Advertising Dept., 345 E. 47th St., New York, N.Y. 10017

## The University of British Columbia Department of Electrical Engineering NSERC/MDA Industrial Research Chair in Radar Remote Sensing and Signal Processing

Applications are invited for a tenured or tenure-track Professor or Associate Professor appointment in Radar Remote Sensing and Signal Processing. The position is funded, initially, under the terms of the Industrial Research Chair program of the Natural Sciences and Engineering Research Council of Canada. The industrial partner, MacDonald Dettwiler (MDA), is providing substantial financial support, with some contributions from the Science Council of British Columbia and the Advanced Systems Institute of British Columbia. Chair funding includes salary support for the Chairholder, support personnel, and graduate students, as well as support for equipment, software, travel and other operating expenses.

A Ph.D. is required. Demonstrated expertise in radar remote sensing and signal processing is essential. Industrial and teaching experience is desirable. The successful applicant will be expected to provide outstanding academic leadership under the terms of the Research Chair Agreement. Substantial collaboration and interaction with the industrial partner is anticipated. Salary, rank and term of appointment will be commensurate with qualifications and experience. The appointment is subject to final budgetary approval.

The University of British Columbia encourages applications from qualified women and minority applicants. In accordance with Canadian Immigration requirements, priority will be given to Canadian citizens and permanent residents of Canada.

The position is available from January 1, 1993. Priority will be given to applications received before November 1, 1992. To apply, send curriculum vitae, reprints of published papers, ■■■ of at least five references, and a statement of eligibility for employment in Canada to:



Dr. R.W. Donaldson, Head  
Department of Electrical Engineering  
The University of British Columbia  
2356 Main Mall, Vancouver, B.C. CANADA V6T 1Z4

## SENIOR

Southern California  
based, non-defense,  
Fortune 500 company

seeks ■ world class Senior Electrical Engineer.

This unique opportunity calls for ■ world-class technical professional with solid analog and digital experience, as well as extensive hands-on

## ELECTRICAL

prototype/project building experience. Possessing a Master's or Doctorate in Electrical Engineering, the ideal individual for this very unique position will be a brilliant "garage/basement tinkerer"; someone with exposure to a wide array of audio, video, telecom-

munications,  
and other  
commercial  
applications.

## ENGINEER

Additionally, you must have demonstrated expertise in the electronics fields, solving electronic problems in circuit design and development.

We offer an outstanding compensation and benefits package. Individuals who meet the above qualifications are invited to submit their resume in strict confidence to our Confidential Reply Service.

P.O. Box 9665, Dept. SEE  
Marina del Rey, CA 90295

Principals Only ■ Equal Opportunity Employer M/F/D/V



# RISC platforms. Innovative operating system. Integrated OLTP solutions. World class database. Open networking solutions. Client server computing. High-speed routing. Meet the challenge of NonStop Availability.

NonStop Availability, the concept of a business-critical OLTP system that never requires any scheduled downtime, forms the heart of Tandem's plans for future product difference. An important piece of our product strategy is providing our customers with the most cost-effective OLTP systems in the industry. Our initiatives in RISC and workstation disk technology are directly in support of our price/performance strategy.

In the area of networking, Tandem's open networking to other computers using SNA, OSI, TCP/IP and other industry-standard protocols is our focus in moving toward open systems. This movement is mirrored in our widening lead in OLTP and database technology.

Tandem seeks a group of senior professionals to complement its elite team of developers. If you're ready to meet the challenge that NonStop Availability provides, apply now to the appropriate address for one of these rewarding positions.

## **NETWORKING/ DATACOMMUNICATIONS**

- Section Managers
  - X.25 and (CAM) Call Application Manager Products
- LAN Communications Platform Development
- LAN Access Methods (UNIX®, Streams) Development
- SNMP Development
- TCP/IP Development and TCP/IP QA
- WAN Access Methods Development and QA
  - 3270 Emulation
  - Bit Synchronous Protocols (HDLC, SDLC, LAPB)
- X.25 Development and X.25 QA
- Hardware Project Design Networking Engineers
- OSI/MAP Development and QA

- X.400 Development and QA
- NLS (National Language Support)/Internationalization Development and QA
- Network Management Development and QA
- SNA Development and QA

Send your resume and salary history to: Tandem Computers, Inc., Attn: Melissa Finder, Dept. SPEC, 10501 North Tantau Avenue, LOC: 201-02, Cupertino, CA 95014-0709, or FAX it to (408) 285-4324.

## **SYSTEM SOFTWARE**

- Operating Systems Kernel Development and SQA
- Remote Diagnostic Services Development
- Systems/Firmware Diagnostics Development
- High level CPU Diagnostic Subsystem Development
- Console Software Quality Assurance Development
- Technical Writing/Hardware/ System Software

## **HARDWARE**

- SCSI Interface/Firmware Engineer
- Peripherals Test/Qualifying Design Engineer
- Storage Products Program Management
- Signal Integrity Design Engineering
- Disk Reliability Engineering
- Component Failure Analysis Engineering
- Software CAD Tools Development
- Software Printer Spooler Development

Send your resume and salary history, to: Tandem Computers, Inc., Attn: Melinda Rowe, Dept. SPEC, 19333 Vallco Parkway, LOC 3-06, Cupertino, CA 95014-2599, or FAX it to (408) 285-7079.

## **OLTP**

### **Operating Systems**

- UNIX Kernel Development
- POSIX™ and X/OPEN Commands & Utilities Development and QA
- File System Development
- System Performance Monitor Development

### **Database Systems**

- SQL Software Development and QA
  - File System, Performance, Catalog Manager, Optimizer, Runtime Executor and Utilities
- Security Systems Development
- DCE Security Services and DCE Directory Services Software and QA Development
- Pathway (TP Monitor) Development and QA
- Low-level Database QA
- Safeguard (Security Product) QA
- TMF QA


### **Software Technical Writers**

- Operating Systems Software and Utilities
- Application Development Tools
- UNIX-based Tools
- National Language Support Software Tools

### **Languages and Tools**

- C Compiler and COBOL85 Development and QA

Send your resume and salary history, to: Tandem Computers, Inc., Attn: Jessi Kaur, Dept. SPEC, 10100 North Tantau Avenue, LOC 251-18, Cupertino, CA 95014-2542, or FAX it to (408) 285-2227. An equal opportunity employer.





## CHAIRPERSON

Department of Electrical Engineering  
State University of New York  
at Stony Brook

The College of Engineering and Applied Sciences at SUNY Stony Brook is conducting a search for a chairperson of the Department of Electrical Engineering effective Spring 1993 or earlier.

The University has selected Electrical Engineering as an area for special emphasis and development. Substantial resources will be made available to the new Chair for this purpose.

Excellent scholarship and leadership qualities are prerequisites for this position. Candidates are expected to have strong research backgrounds as indicated by publication history and status within his or her field. Of special importance will be the chairperson's willingness and ability to work closely with a diverse community of scholars, industrial corporations, and government agencies.

The Department, one of six in the College of Engineering and Applied Sciences, has 20 full-time faculty. An additional 5 faculty members have joint appointments with other departments. Six of the faculty are Fellows of IEEE. Many have substantial professional activity and have gained other major recognition. The Department offers B.E., M.S. and Ph.D. degrees. There are currently 380 undergraduates enrolled in two programs—Electrical Engineering and Computer Engineering. Both programs are accredited by ABET. Graduate student enrollment includes 94 full-time students. Faculty research interests include Bioengineering, Computer-Aided Design, Communication Systems and Networks, Computer Architecture, Control Systems, Fault Tolerant Computing, Machine Vision, Microwave Electronics, Optoelectronics, Signal Processing, and Analog/Digital VLSI.

SUNY Stony Brook is a major research university in the State of New York located on the scenic north shore of Long Island in a picturesque community 60 miles east of New York City. Campus student enrollment is 17,600 and there are 1,540 full-time faculty members. The College of Engineering and Applied Sciences is one of five major academic units which altogether offer graduate degrees in more than 43 subject areas, most of them including doctoral studies.

Applications and nominations should be addressed to: Professor Phillip M. Lewis, Chair, Electrical Engineering Search Committee, Computer Science Department, SUNY Stony Brook, Stony Brook, NY 11794-4400. Applications should include a vita and the names of at least three references. SUNY Stony Brook is an affirmative action/equal opportunity educator and employer. AK173

## VICE PRESIDENT AND DIRECTOR GEORGIA TECH RESEARCH INSTITUTE

Nominations and applications are invited for the position of Vice President and Director of the Georgia Tech Research Institute (GTRI). The GTRI Vice President reports to the Executive Vice President of Georgia Tech.

GTRI is a nonprofit, client-oriented organization with significant and expanding interactions with Georgia Tech's academic faculty and students. The GTRI Director will be responsible for the technical and administrative leadership of GTRI, enhancing the image of GTRI with outside funding sources, developing new funding sources, expanding Georgia Tech's research and development capabilities and promoting increased integration with Georgia Tech's academic programs. This position will be responsible for the management and supervision of a staff of approximately 1,500, which includes more than 600 full-time researchers, and a research program of over \$100 million annually. GTRI is internationally recognized for its expertise in areas such as aerospace science and technology, artificial intelligence, computer science and technology, electronic systems, electronic warfare, environmental science, infrared and electro-optics, manufacturing research, materials science, radar, signature control technology, and threat systems development.

The mission of GTRI includes public service, and outreach and technology transfer across a broad spectrum of scientific and technological endeavors in addition to research and development.

Candidates for the position should hold a degree in Engineering or Science, with a strong preference for a doctoral degree. The successful candidate should have a distinguished record of management and development of research, strategic planning, human resource development and Total Quality Management.

Nominations and applications should be sent to:

**Dr. Gary Poehlein, Office of Interdisciplinary Programs, Code IES83,  
Georgia Institute of Technology, 400 Tenth Street, Centennial Research  
Building, Room 282, Atlanta, GA 30332-0370, Telephone (404) 894-4826.**

Review of applications will begin as soon as they are received and continue until the position is filled. Georgia Institute of Technology is a part of the University System of Georgia and is an Equal Opportunity Employer.

*Georgia Tech*

## FRONTIERS: ALASKA AND SPACE

A well-established research institute is looking for an innovative person to manage an exciting earth-observing satellite project.

### PROJECT MANAGER, ALASKA SAR FACILITY GEOPHYSICAL INSTITUTE UNIVERSITY OF ALASKA FAIRBANKS

The Geophysical Institute of the University of Alaska Fairbanks is seeking a project manager for the Alaska SAR Facility, a joint NASA-University of Alaska facility receiving, processing, and analyzing synthetic aperture radar data from the European ERS-1, Japanese JERS-1, and Canadian Radarsat satellites. The Alaska SAR Facility is one of seven DAACs (Distributed Active Archive Center) of NASA's EOS, the Earth Observing System program. The project employs about 35 people and has an annual budget of over three million dollars.

The successful candidate must have considerable experience in the establishment, operation, and management of satellite ground stations, or in major projects related to such ground stations. Expertise in the technical, personnel, and financial aspects of managing large projects of this nature are essential. An advanced degree in engineering or science is required.

Duties include technical supervision of the facility; providing an interface with NASA, contractors, and foreign satellite flight agencies and ground stations; planning and implementing facility upgrades; hiring and review of the technical personnel; production of deliverables; and budgetary responsibility.

Salary will depend on qualifications and experience. The university has an excellent retirement and benefits program. Please reply by 1 November 1992 to the address below, or ask for additional information:

Director, Alaska SAR Facility  
Geophysical Institute  
University of Alaska Fairbanks  
Fairbanks, AK 99775-0800  
Fax. (907) 474-5195  
Tel. (907) 474-7371



The University of Alaska is an EO/AA employer and educational institution. Applications from minorities and women are especially encouraged.



## Recent books

(Continued from p. 78)

**Elements of Project Management: Plan, Schedule & Control, 2nd edition.** *Spinner, M. Pete*, Prentice Hall, Englewood Cliffs, N.J., 1992, 211 pp., \$40.

**Using MS-DOS Kermit: Connecting Your PC to the Electronic World, 2nd edition.** *Gianone, Christine M.*, Digital Press, Bedford, Mass., 1992, 345 pp., \$34.95.

**DB2: Performance, Design, and Implementation.** *Silverberg, David*, McGraw-Hill, New York, 1992, 381 pp., \$44.95.

**RISC/CISC Development and Test Support.** *Hobbs, Marvin*, Prentice Hall, Englewood Cliffs, N.J., 1992, 386 pp., \$46.

**McGraw-Hill Dictionary of Information Technology & Computer Acronyms, Initials & Abbreviations.** *Rosenberg, Jerry M.*, McGraw-Hill, New York, 1992, 209 pp., \$24.95.

**Neural Networks for Signal Processing.** *Kosko, Bart*, Prentice Hall, Englewood Cliffs, N.J., 1992, 399 pp., \$50.

**McGraw-Hill Dictionary of Wall Street Acronyms, Initials & Abbreviations.** *Rosenberg, Jerry M.*, McGraw-Hill, New York, 1992, 235 pp., \$24.95.

**DECnet Phase V: An OSI Implementation.** *Martin, James*, and *Leben, Joe*, Digital Press, Bedford, Mass., 1992, 590 pp., \$49.95.

**McGraw-Hill Dictionary of Business Acronyms, Initials & Abbreviations.** *Rosenberg, Jerry M.*, McGraw-Hill, New York, 1992, 352 pp., \$29.95.

**VAX/VMS Operating System Concepts.** *Miller, David Donald*, Digital Press, Bedford, Mass., 1992, 512 pp., \$44.95.

**The Entrepreneurial PC.** *David, Bernard J.*, Windcrest/McGraw-Hill, Blue Ridge Summit, Pa., 1992, 319 pp., \$19.95.

**Developing International User Information.** *Jones, Scott, et al.*, Digital Press, Bedford, Mass., 1992, 214 pp., \$24.95.

**The Laser Guidebook, 2nd edition.** *Hecht, Jeff*, McGraw-Hill, New York, 1992, 498 pp., \$44.95.

**Electric Circuit Analysis, 2nd edition.** *Johnson, David E.*, Prentice Hall, Englewood Cliffs, N.J., 1992, 779 pp., \$61.

**The Stephen Cobb Complete Book of PC and LAN Security.** *Cobb, Stephen*, Windcrest/McGraw-Hill, Blue Ridge Summit, Pa., 1992, 556 pp., \$24.95.

**The Computer Professional's Survival Guide.** *Simon, Alan R.*, McGraw-Hill, New York, 1992, 176 pp., \$24.95.

**Control Systems Engineering.** *Nise, Norman S.*, Benjamin/Cummings Publishing, Redwood City, Calif., 1992, 756 pp., \$64.50.

**The Art of Technical Documentation.** *Haramundanis, Katherine*, Digital Press, Bedford, Mass., 1992, 272 pp., \$28.95.

**Introduction to Parallel Algorithms and Architectures: Arrays, Trees, Hypercubes.** *Leight-*

*on, F. Thomson*, Morgan Kaufmann Publishers, San Mateo, Calif., 1992, 831 pp., \$54.95.

**Power System Transient Stability Analysis: Using the Transient Energy Function Method.** *Fouad, A. A.*, and *Vittal, Vijay*, Prentice Hall, Englewood Cliffs, N.J., 1992, 357 pp., \$65.

**Profiting from Innovation: The Report of the Three-Year Study from the National Academy of Engineering.** Eds. *Howard, William G., Jr.*, and *Guile, Bruce R.*, Free Press/Macmillan, New York, 1992, 154 pp., \$22.95.



## CREATE THE EIGHTH WONDER OF THE WORLD.

Only the most impressive artistic and engineering achievements of the ancient world came to be known as "wonders." Today at NOKIA Mobile Phones in San Diego, our professionals are engineering modern wonders that are making history with their progressive capabilities and features.

Finnish-based NOKIA Mobile Phones is Europe's largest—and the world's second largest, mobile phone manufacturer. As an Engineer with our new U.S. Research and Development Center in San Diego, you can use your creativity and skills to create NOKIA's next-generation digital cellular mobile phones. You will be supported by our progressive operational and management style—one that has earned us a very high rating for TQM, and has led to the implementation of concurrent engineering. And you'll make history with the contributions you make to our product line.

Our current openings include:

**HARDWARE ENGINEER** Requires a BS/MS in Electrical Engineering or related discipline, and 5+ years' industry experience. Familiarity with CAD (Mentor, PCAD, Orcad), project and test planning, EMI and Product Qualifications, mixed technology and analog design required. Experience with entire product cycle of a self-contained product is a must.

**PROJECT LEADER** Requires a BS, with MS or Ph.D preferred in Electrical Engineering, and 8-10 years' engineering experience related to digital wireless communications products. Proven project leadership in developing and implementing commercial cellular telephone products is strongly preferred.

**IC DESIGN ENGINEER** Requires a BS/MS/Ph.D in a related discipline and 5 years' experience in analog CMOS or Bi-CMOS integrated circuit design. Must be familiar with amplifiers, A/D, and D/A converters and low power design. Knowledge and ability to use modern tools in Apollo/HP environment are necessary.

Help us create wondrous achievements for our industry—and our time—by joining NOKIA Mobile Phones. For additional information and other career opportunities available, send your resume to: NOKIA Mobile Phones, Code 121, 9605 Scranton Rd., #303, San Diego, CA 92122, or fax it to (619) 450-3168. Principals only.





### RESPONSIBILITIES:

Provide design expertise for development of advanced high-performance free-electron laser systems. Responsible for conceptualizing, designing, building, and modifying particle accelerator electrical and electronic subsystems. Subsystems include high pulsed power microwave generations, high-voltage systems, electron beam transport systems, personnel safety and hardware protection systems, and beam diagnostics and instrumentation. Responsible for the implementation, operation, and maintenance of a computer-based control and interlock system. Participate in the design and execution of experiments. Responsible for documentation of experimental test procedures and results.

### EXPERIENCE REQUIREMENTS:

Demonstrated experience and understanding of electrical and electronics systems associated with particle accelerators. Demonstrated expertise in analog and digital circuit design and implementation. Demonstrated ability to design and implement fiber optic and EMI noise suppression systems. Experience in the design of particle beam diagnostics instrumentation. Demonstrated experience in micro-processor programming. Effective oral and written communication skills must be documented. Demonstrated rigorous experimental technique through effective and thorough documentation of experimental procedures and test data.

### EDUCATIONAL REQUIREMENTS:

Associates degree in Electronics Engineering and demonstrated accelerator experience or equivalent formal training plus demonstrated accelerator experience.

#### Definition of Equivalent Training:

Equivalent formal training in level and purpose could include: (1) a 70-hour program of study in Electronics Engineering Technology from an accredited Junior College; (2) successful completion of the first three years of a University Bachelor's-Degree Program in Electronics Engineering; or (3) a diploma from a Community College resulting from a three-year program of study in Electronics Engineering Technology.

**SALARY:** \$44,554.00

#### Send resumes to:

New Mexico Department of Labor  
P.O. Box 1908  
Santa Fe, NM 87501

Job Order No. 334067

Affirmative Action/Equal Opportunity Employer

### 44Hr. College Credit Course

### CAD for Printed Circuit Board Design using

# TANGO

also available

**New York State  
Professional Engineer's  
Review Course**

**Starts Sept. 12, 1992**

**Call Prof. Rus Pyros  
TECHNICAL CAREER  
INSTITUTES**

**212 - 594 - 4000**

Circle No. 3

## Software Platforms Systems Developers/Systems Engineers

Come join a \$380 million worldwide provider of information management solutions. **Cincinnati Bell Information Systems (CBIS)** is seeking individuals with a solid background in C/UNIX to apply an open-system platforms approach to develop new **CBIS Edge<sup>sm</sup>** software products for our clients.

The new **CBIS Edge** represents an area of unlimited personal and professional growth opportunities for select individuals in our Software Development Centers in Cincinnati, suburban Chicago, and Orlando. Competitive systems development and systems engineering candidates will have three or more years' experience in one or more of the following: C++, Object-Oriented Programming, Open-Systems Architecture, Case Tools, and Requirements Definition and Analysis.

**CBIS** provides merit-based compensation plans, relocation assistance, "FLEX" benefits, 401(k) savings plan, and tuition reimbursement. Please mail your resume to **CBIS, P.O. Box 1638, Dept. CD392, Cincinnati, OH 45201, ATTN: Steve Suiter, Director of Organizational Development.**



We are an Equal Opportunity/Affirmative Action Employer.  
Edge is a service mark of CBIS.

### CONTROL SYSTEMS & ELECTRICAL Engineers

GILBERT/COMMONWEALTH, INC. has new opportunities for engineering professionals at our **Reading-PA** Hqs., in our **Oak Ridge-TN** SE regional hqs., and in other U.S. locations. Engineering degree essential. Openings include:

**FOSSIL POWER PLANT:** Min. 4 yrs. exp. in engrg. of DCS's and instrumentation or electrical equipment & power distribution systems for plant modernization projects required.

**INDUSTRIAL:** 2-10 yrs. exp. in engrg. & development of process control and manufacturing automation systems, or modernization & expansion of electrical systems & equipment for process & manufacturing industries required.

Excellent interpersonal and communication skills and a strong client orientation are essential; P.E. license a plus. Both full-time and hourly positions are available. Full-time positions qualify for benefits, our relocation program, and the option to work a "9/80" schedule and have EVERY OTHER FRIDAY OFF! Send resume to **D. Scibek, Human Resources/S992, P.O. Box 1498, Reading, PA 19603-1498. FAX: (215)775-2670.** Equal Opportunity Employer  
M/F/D/V. An ADNET ONLINE Subscriber.



**Gilbert/Commonwealth, Inc.**  
Engineers and Consultants



# Explore Opportunities with CAE Leader!

## Development Engineer

As a member of our R&D staff, you will participate in the creation of sophisticated CAE tools. Particular areas of involvement will include development of time- and frequency-domain simulation algorithms for high-frequency circuits and active/passive element models for inclusion in EEsof's CAE products. Requires: MSEE (Ph.D. preferred); 2+ yrs. exp. in CAE as applied to linear/nonlinear circuit analysis in time- and frequency-domain, SPICE-type simulation algorithms, and high-frequency active/passive element models.

## Senior Scientist

As part of our Product Marketing team, you will serve as a senior resource for RF systems theory and design. Will also provide technical liaison between customers, Marketing, and R&D on systems simulation tools. Requires: MSEE (Ph.D. preferred); 5+ yrs. exp. in system theory and design (with emphasis in digital communications); detailed theoretical understanding of communications, EW, and RADAR systems; and strong writing and presentation skills. System simulation experience highly desirable.

As the leading developer of high-frequency analog and microwave EDA tools, EEsof offers competitive compensation and full benefits in a technically challenging environment. Explore opportunities by mailing or faxing your resume (principals only), in confidence, to: EEsof, Attn: HR-IEEE 992, 5601 Lindero Canyon Road, Westlake Village, CA 91362. FAX: 818-879-6212.

Equal Opportunity/Affirmative Action Employer

**EEsof** Software Solutions for  
High-Frequency Analog Design

## SYSTEMS ENGINEER



## Ground Vehicle Electronic Systems

A division of the Fortune 300 Harsco Corporation and a major producer of military vehicles, we have a worldwide reputation for excellence and innovation. Reporting to the Chief Staff Engineer, this newly created position will be responsible for performing requirements analysis and allocation for combat vehicles.

A BSEE or equivalent background and 5+ years' military ground vehicle electronic systems research and development experience are required. Knowledge of system/subsystem nuclear survivability is desirable.

Our Southcentral PA location just one hour North of Baltimore offers quality, affordable living with recreational, healthcare and educational opportunities to meet your needs. We offer competitive compensation and benefits. Please send your resume to: **Human Resources/S992, P.O. Box 15512, York, PA 17405-1512.** Direct responses only, please. BMY is an equal opportunity employer.

**BMY** Combat  
Systems  
A DIVISION OF **harsco**  
CORPORATION

# INFINITE CHALLENGES

E-Systems ECI Division is in need of engineers with military satellite communication experience (ground, manpack/man-portable, airborne, or spaceborne); a BSEE/BSCE; and at least 2 years' experience in one of the following areas:

## SOFTWARE

- ADA, C
- 1750A, of 68020 Microprocessors
- VAX VMS, Sun UNIX
- Real-time, Embedded Microprocessor
- DOD-STD-2167A, CASE Tool

## DIGITAL HARDWARE

- ACTEL FPGAs
- Microprocessor based systems
- 1553 bus interface

## DIGITAL SIGNAL PROCESSING

- Discrete Fourier Transforms
- Control Loops
- PSK Demodulation

## EMBEDDED CRYPTO

- Security Fault Analysis
- TEMPEST, Red/Black Isolation
- Related interface hardware

## RF and MICROWAVE

- Synthesizer Design, Direct Digital
- Power amp and filter design
- MMIC design

## ANTENNA DESIGN

- Parabolic Antenna Design
- Gimbal, Positioner
- 10 TO 60 GHz

## SYSTEMS

- BSEE/MSEE, minimum 4 years' experience
- Strong communication background
- Requirements Analysis, Functional Analysis
- System Synthesis, System Analysis
- RF Link Budget Analysis
- System Integration/Test
- Customer Interface

Background: Hardware, Software Architecture Cryptographic; BIT/BITE; Antenna Pointing, Tracking, and Platform Stabilization; MIL-STD-1582; SI-1135, SI-2035, LL1005.

MILSTAR or other military satellite design experience a plus.

E-Systems offers very competitive salaries and an excellent benefits package which includes an Employee Stock Ownership Plan, 401 (k), and major medical and dental insurance. Qualified candidates should forward a resume and salary history to: Manager of Staffing, E-Systems, Inc., ECI Division, Post Office Box 12248, St. Petersburg, Florida 33733-2248.



**E-SYSTEMS**  
The science of systems.

U.S. Citizenship Required.  
An Equal Opportunity Employer, M/F, D/V.





**ONLY A  
FEW PEOPLE  
WILL KNOW  
WHAT YOU'RE  
WORKING ON.**

When you're involved with the most advanced, sensitive areas in aviation defense technology, your current accomplishments are recognized by just a handful of people.

But think who those people are.

For over 40 years, we've been counted on to deliver breakthroughs vital to our nation's security. We're Lockheed Advanced Development Company (LADC), the legendary Skunk Works. We've built a reputation for doing the impossible on time and within budget.

A major component of LADC is our Electromagnetics Division, which is engaged in a variety of programs with advanced performance requirements including low observability. To help us in this critical work, we need Antenna/Radome Designers and Anechoic Range Test Engineers who have expertise in one or more of the following areas:

Analysis and design of antennas and Radomes:

- DF arrays
- Phased arrays
- Conformal antennas
- Radomes and frequency selective surfaces

Analysis and design of microwave components:

- Stripline circuitry
- Power distribution networks
- RF filters
- Packaging and mechanical design of microwave components, antennas and radomes

Electromagnetic Testing, including Anechoic Range Operations:

- RCS studies
- Nearfield measurements
- Antenna/Radome evaluation

Applicants selected will be subject to a security investigation and must meet eligibility requirements for access to classified information. Our company offers competitive salaries, based on experience, and a generous benefits package. For immediate consideration, please send your resume to: Lockheed Advanced Development Company, Professional Staffing, Dept. 0041, P.O. Box 250, Sunland, CA 91041.

*Lockheed is an equal opportunity employer m/f/h.*

 **Lockheed**  
Advanced Development Company




**INNOVATION: giving shape to imagination.**

SKUNK WORKS and the skunk design are registered service marks of the Lockheed Corporation. (c)1992 Lockheed Corporation



# MICRO SYSTEM Technologies 92

3<sup>rd</sup> Int. Conference and Technical Exhibits on Micro-Electro, Opto, Mechanical Systems and Components



## Come to Berlin!

### Meet the key persons in micro-machining:

Design, manufacturing and application of micro-structures. The integration of micro-mechanics, micro-electronics, micro-optics and micro-sensing. The meeting for specialists and users worldwide. Applications in computer and communication technologies, consumer electronics, automotive electronics, environment engineering, medicine technologies...



IEEE Institute of Electrical and Electronic Engineers

MSTA World Micro Systems Technology Association



International Congress Center Berlin  
21 - 23. Okt. 1992

### Conference Themes (excerpt)

- Microrobotics - Approach to the Realisation
  - International Support Programs
  - Bionik, Biosensors
  - High Temperature Superconductors
  - Packaging
  - Materials
  - Sensor/Actuator Principles
  - Micro Optics
  - Optical Methods for Characterization
  - Micromachining
  - Future Medical Applications
  - Business Opportunities
  - Application in Consumer Products
  - Gas Sensors/Chemical Sensors
  - Silicon Technologies & Integration
  - Automotive Applications
  - Micro System Design
  - Advances in Plotter Technologies
  - Measurement Application/Control
  - Device Simulation
  - Telecommunication Applications
- Featuring the key speakers from all over the world.  
Conference language is English.

### Organizers:

AMK Berlin Ausstellungs-Messe- und Kongreß GmbH  
MESAGO Messe & Kongreß GmbH  
With the support of the Senate of Berlin, Department of Economics

### Book your participation now

Simply copy this part of the ad and fax it with your name +49-711-6194697

### Information:

ICC BERLIN  
Hallerstrasse 1  
P.O. Box 1885  
10000 Berlin, Germany  
Tel. (030) 3836711  
Fax (030) 3836711

Circle No. 17

## 0.1 to 20 GHz Circuit Design/Antenna Design

Join our research department in Murray Hill, NJ to pioneer technologies applicable to wireless computing/communications products. We are seeking an outstanding research leader to explore new concepts for antenna design, low power r.f. and digital circuits and power management. This position requires a graduate degree (Ph.D. preferred) in Electronic Engineering or Physics.

Your previous experience should include one or more of the following: r.f. circuit design including development of components for filters and antennas; design of low power circuits in GaAs or silicon; simulation and analysis (including hands-on testing) of circuit performance from 100 MHz to 20 GHz. Prior experience on product development and/or manufacturing of related products is also desirable.

You will possess an ability to identify and formulate effective solutions to unstructured problems and successfully accomplish their implementation. You must also have demonstrated creativity through patents, original technical publications or innovative product development and will be capable of working with product developers/manufacturers/marketers in business units. Outstanding interpersonal, organizational and communication skills are essential.

For consideration, send your resume to: AT&T Bell Laboratories, Technical Employment Manager, Dept 86/6331/92, P.O. Box 636, 600 Mountain Avenue, Murray Hill, NJ 07974-0636. An equal opportunity employer.



**AT&T**

The right choice.

**MIT**

## FACULTY SEARCH MATERIALS SCIENCE MATERIALS ENGINEERING

The MIT Department of Materials Science and Engineering has searches in progress for tenure track Assistant Professors.

Faculty members are sought with demonstrated potential, creativity and teaching ability in all aspect of the field of materials science and engineering. These positions are potentially open for persons specializing in any of the materials classes (metals, ceramics, polymers or electronic materials), or more broadly in materials science or materials engineering.

Doctorate required and experience desirable. Salaries and benefits competitive. Positions will be filled only if exceptional candidates are found. Appointment to a more senior position is possible based on experience and accomplishments.

Interested persons should send a complete resume including a list of publications and names and addresses of three persons who could supply letters of reference. This material should be sent to: Professor Merton C. Flemings, Head, Department of Materials Science and Engineering, Massachusetts Institute of Technology, Building 8-309, 77 Massachusetts Avenue, Cambridge, MA 02139-4307.

MIT is an Equal Opportunity/Affirmative Action Employer  
MIT is a non-smoking environment

Massachusetts  
Institute of Technology



## CLASSIFIED EMPLOYMENT OPPORTUNITIES

The following listings of interest to IEEE members have been placed by educational, government, and industrial organizations as well as by individuals seeking positions. To respond, apply in writing to the address given or to the box number listed in **■** of *Spectrum Magazine*, Classified Employment Opportunities Department, 345 E. 47th St., New York, N.Y. 10017.

### ADVERTISING RATES

**Positions open**—\$36.00 per line, not agency-commissionable

**Positions wanted**—\$36.00 per line, **■** 50% discount for IEEE members who supply their membership numbers with advertising copy

All classified advertising copy must be received by the 25th of month, two months preceding date of issue. No telephone orders accepted. For further information contact Francesca Silvestri, 212-705-7578.

IEEE encourages employers to offer salaries that **■** competitive, but occasionally **■** salary may be offered that is significantly below currently acceptable levels. In such cases the reader may wish to inquire of the employer whether extenuating circumstances apply.

### Academic Positions Open

**Position Open:** The Dept. of BME of the University of Miami seeks applications and nominations for a new tenure track faculty position at the Asst. Prof. rank, for the Fall '92 semester, with the following minimum qualifications: PhD in Engineering or in a closely allied field; research experience in BME; demonstrable interest in the development and teaching of undergraduate level lecture and lab courses in BME. The 9-month appointment may be extended for summers. Salary and fringe benefits **■** competitive. The University of Miami is an equal opportunity employer. Please contact: Dr. Ozcan Ozdamar, Dept. of BME, College of Engineering, University of Miami, POB 248294, Coral Gables, FL 33124.

**Yale University**—Electrical Engineering. The Department of Electrical Engineering invites applications for a faculty position at the Assistant or Term Associate Professorial level in VLSI system design, computer architecture, parallel processors, and other areas of computer engineering. Applicants should have a Ph.D. in Electrical Engineering, Computer Science, or closely related field, and should exhibit outstanding research accomplishments and **■** commitment to teaching. Close collaboration with the existing computer engineering group in the Department and interaction with the Department of Computer Science **■** desired. Preferred candidates should also have **■** strong interest in other programs in the Department, including microelectronics, systems science, and signal processing, and would be expected to contribute to joint research activities. Applications are accepted until October 30, 1992. Please send resume, including names, addresses, and telephone numbers of at least three references to: Professor T.P. Ma, Chair, Department of Electrical Engineering, Yale University, 15 Prospect Street, New Haven, CT 06520-2157. Yale University is an affirmative action, equal opportunity employer.

**Director**—The School of Electrical Engineering and Computer Science at Washington State University is soliciting applications and nominations for the position of Director of the School. The School presently has 39 full-time faculty, 452 undergraduates, and 171 graduate students. It awards B.S., M.S., and Ph.D. degrees in electrical engineering and computer science. Since 1985 the annual research budget has quintupled to \$1.6 million. Major **■** of re-

search interest include energy systems, solid state electronics, communications and signal processing, electromagnetics, control theory, software engineering, graphics and visualization, and parallel computing. The School enjoys strong ties with local and national industries in **■** number of educational and research areas. Recently, the Boeing Company established the Boeing Centennial Chair in Computer Engineering, and the School has distinguished professorships in electromagnetics, analog electronics, and power engineering. In addition, an NSF Center for Design of Analog-Digital Integrated Circuits, with 15 industrial members and 3 affiliated universities, is based in the School. While our faculty represent diverse research interests, we also are united in our commitment to education. In recent years, several faculty members have been awarded NSF grants under the ILI Program and several have published textbooks. In the past year the industrial community has donated \$3.1 million in state-of-the-art laboratory equipment and software to help maintain our tradition of "hands-on" education. Candidates must (1) possess a Ph.D. degree in electrical engineering, computer science, or **■** related discipline; (2) demonstrate outstanding leadership, an international research reputation, and a clear vision of the growth and development needed to enhance excellence in both teaching and research; (3) possess good communication skills and the ability to motivate people; (4) be able to interact productively with universities, industry, and government; (5) be committed to the promotion of faculty research activities and able to mentor junior faculty members to assist them in establishing strong research programs; (6) have **■** commitment to excellence in undergraduate education; and (7) be able to take the lead in the continued development of **■** diverse work force. Washington State University is a land grant university located in the beautiful rolling hills of Eastern Washington. The City of Pullman offers an ideal setting for family life; it has one of the best public school systems in the state, and recreational and outdoor activities **■** available year-round. The State of Washington has no income tax, real estate taxes **■** reasonable, and housing costs in the Pullman **■** moderate. Applications and nominations should be sent to EECS Director Search Committee, School of Electrical Engineering and Computer Science, Washington State University, Pullman, WA 99164-2752. Screening of applicants will begin Nov. 1, 1992, and will continue until the position is filled. Washington State University is an equal opportunity/affirmative action employer. Members of ethnic minorities, women, Vietnam **■** or disabled veterans, persons of disability and/or persons between the ages of 40 and 70 are encouraged to apply.

**Deane Ackers Professor** of Electrical and Computer Engineering. The University of Kansas Department of Electrical and Computer Engineering invites applications and nominations for a chaired professor. Appointment will be at the rank of tenured full professor. The Deane Ackers Professor of Electrical and Computer Engineering will be expected to teach and lead a sponsored research program in microwave remote sensing and electromagnetics. Candidates must have an earned Ph.D. in electrical engineering, a demonstrated record of teaching both undergraduate and graduate courses, **■** outstanding record of research in microwave remote sensing and electromagnetics, and an international reputation. The Radar Systems and Remote Sensing Laboratory (RSL) has **■** long history of graduate education and research. Through RSL, the University of Kansas is internationally recognized for its training and research activities in the fields of radar remote sensing and applied remote sensing and, based on the most recent survey, granted more remote sensing doctorates than any other university in the country between 1965 and 1984. In addition to research in remote sensing, the Electrical and Computer Engineering Department has active research programs in telecommunications and information sciences, and computer engineering. The successful candidate is expected

ed to lead the Radar Systems and Remote Sensing Laboratory. He/she should be able to provide leadership in research, to teach effectively, and to establish significant collaborations with governmental and industrial researchers and engineers. Salary will be commensurate with the outstanding qualifications required of the appointee. Applications and nominations with a resume and a list of at least three references should be sent to Professor Prasad Gogineni, chairman of the search committee, Electrical and Computer Engineering Department, 1013 Learned Hall, Lawrence, KS 66045. Review of applications will begin on July 15, 1992 and continue until the position is filled. The University of Kansas is an affirmative action/equal opportunity employer.

**The Institute for Micromanufacturing** at Louisiana Tech University is inviting applications from qualified individuals for tenure track faculty, research engineers, research associates, visiting scholars, postdoctors and technician positions in the **■** of optical lithography, x-ray lithography, metrology, materials science, micromechanics, and technology transfer to support the development of microelectromechanical systems (MEMS) to the prototype stage. Appointments for the faculty positions will be considered at the Assistant, Associate and Full Professor rank commensurate with qualifications, which include **■** earned doctorate in mechanical engineering, electrical engineering, materials science, physics, chemistry, biology or **■** related field, and a strong commitment to education, developing externally funded research. Group leaders, qualified at the full professor level, are being sought in several areas. The research engineer and associate positions require a minimum of **■** BS degree and appropriate professional experience in one of the areas mentioned above. Screening of applicants will begin immediately and applications will be accepted until the positions are filled. Some positions will not be filled until September of 1993 or later. Please send resume, names of three professional references and **■** brief description of teaching and research interests to Dr. Robert O. Warrington, Director, Institute for Micromanufacturing, Louisiana Tech University, P.O. Box 10348, Ruston, LA 71272-0046. Louisiana Tech University is **■** equal opportunity employer. Women and minorities are encouraged to apply.

**E.E. to support university laboratory** dedicated to speech communication/signal processing via audio, vibration, and vision. B.S.E.E. acceptable, M.S. desirable. Reference letters and ability to learn sign language required. Send cover letter and resume to Gallaudet Univ., Personnel Office, 800 Florida Ave., NE, Washington, DC 20002. Equal Opportunity/Affirmative Action Institution.

**Chair in Photonic Systems.** The Department of Electrical Engineering of McGill University is seeking an incumbent for a new Chair in Photonic Systems. Bell Northern Research and Northern Telecom Ltd., together with McGill University, will be seeking support for this Chair from the Natural Sciences and Engineering Research Council (NSERC) of Canada, through their Industrial Research Chair (IRC) program. Since this program is based on peer review, the successful candidate must be an internationally recognized authority, preferably working at the interface between photonic devices and systems, with outstanding scientific and leadership qualities that can bring the Department of Electrical Engineering to the forefront of research in Photonic Systems. An important condition for achieving this goal is the appointee's ability to build up a strong academic program for both graduate and undergraduate students. The position carries a highly competitive salary, junior faculty positions specifically designated in Photonic systems, necessary laboratory space and, above all, strong University support for the above stated goals. Interested candidates **■** invited to send their resume to: Professor N.C. Rumin, Chairman, Department of Electrical Engineering, McGill Univer-



sity, 3480 University St., Montreal, QC, Canada, H3A 2A7. McGill University is committed to equity in employment.

**The Johns Hopkins University**, Department of Electrical and Computer Engineering, invites applications for tenure-track faculty positions at the assistant or associate professor level in the areas of computer engineering; solid state and quantum electronics; and signal and image processing. Candidates for associate professor appointments are expected to have significant research records. Candidates for assistant professor appointments are expected to show strong research potential. Applicants should send resumes, including names of at least three references, to Search Committee, Department of Electrical and Computer Engineering, The Johns Hopkins University, Baltimore MD 21218-2868. The Johns Hopkins University is an equal opportunity/affirmative action employer.

**Faculty Position in Electronics.** University of North Texas - Department of Engineering Technology. The Department of Engineering Technology announces a faculty position commencing August 15, 1993. A doctorate is required for a tenure-track appointment. The minimum academic credential is a masters degree in nuclear or electrical engineering/engineering technology or a closely related field. The successful applicant must have demonstrated competency in nuclear and electrical technology disciplines. Requirements include directing the nuclear technology program and teaching within both the nuclear and electrical engineering technology curricula. Three years of relevant industrial experience and a record of successful teaching are required. Candidates from industry, engineering, and engineering technology programs are encouraged to apply. We are women and minority candidates. In addition to teaching, faculty are expected to engage in program development, advise students, serve on committees, pursue scholarly activity, publish, and participate in professional societies. The University of North Texas is an equal opportunity/affirmative action employer. Submit letter of application, official transcripts, vita, references, and three current letters of recommendation to: Dr. Phillip R. Foster, Search Committee Chair, University of North Texas, Department of Engineering Technology, P.O. Box 13198, Denton, TX 76203, Ph. (817) 565-2022. Applications accepted until position is filled.

**Research Position in Applied Ocean Science**, University of California, San Diego. The Marine Physical Laboratory, Scripps Institution of Oceanography, invites applications for scientists for appointments at the Assistant Research level. Applicants with an interest in conducting innovative experimental work at sea will be given strongest consideration. Fields of interest include, but not limited to, sonar, ocean acoustics, geology, geophysics, geodesy, physical oceanography, submersible vehicles, radar oceanography, and marine micrometeorology. Applicants should have a Ph.D. and will be expected to show evidence of their potential through letters of recommendation and a publication record appropriate for their experience. Salary will be commensurate with experience and qualifications and based on UC pay scale. Immigration status of non-US citizens should be stated in the resume. Closing date for applications is September 21, 1992. Direct inquiries to Dr. K.M. Watson, Acting Director, Marine Physical Laboratory, Scripps Institution of Oceanography, San Diego, CA 92152. (619) 534-1803. The University of California, San Diego, is an equal opportunity/affirmative action employer.

**Brown University**—Faculty Position in Electrical Engineering, Semiconductor Electronics or Optoelectronics. The Division of Engineering at Brown University announces the opening of a tenure track assistant professor position in Electrical Engineering, expected to be filled by July 1, 1993. Applications are invited from candidates who have a Ph.D. or equivalent degree in Electrical Engineering, Applied Physics, or Physics, and who possess a demonstrated record of accomplishment in experimental research with electronic or optoelectronic devices. Areas of preferred specialty include applied and basic semiconductor research such as quantum transport devices, semiconductor

lasers, and nonlinear optics, especially in the context of submicron structures and novel materials. The position is part of an active program in electronic materials and devices at Brown which includes major new facilities. The appointee will be expected to teach undergraduate and graduate courses in the appropriate specialty in Electrical Engineering and in the Engineering Core Curriculum, and to conduct significant, independent research. For full consideration, please respond by November 15, 1992. The search will remain open until the position is filled. Complete resumes, including the names of at least three references should be sent to Professor Maurice Glicksman, Search Committee Co-Chair, Division of Engineering, Box D, Brown University, Providence, RI 02912. Brown University is an affirmative action/equal opportunity employer.

**Seoul National University:** Department of Control and Instrumentation Engineering expects some faculty position openings in the near future and is seeking outstanding faculty candidates with research in the general areas of Control and Instrumentation. Korean citizenship is required. Applicants with strong commitment and outstanding qualification should send a letter of interest with a resume including undergraduate and graduate GPA's, complete publication list, dissertation abstract, post-doctoral experience, and the names of three references to: Dr. In-Joong Ha, Chairman, Department of Control and Instrumentation Eng., Seoul National University, San 56-1 Shinrim-Dong, Kwanak-Ku, Seoul 151-742, Korea.

**The University of Saskatchewan** invites applications for a tenure track position in the Department of Electrical Engineering in the area of Power System Control. The responsibilities include teaching undergraduate courses in electrical engineering and power system control, graduate courses in power system control, and conducting research. Appointments are normally made at the Assistant Professor level. Applicant must hold an earned Ph.D. degree and have demonstrated potential for teaching at the undergraduate and graduate levels and for developing an independent research program. The department offers programs leading to B.E., M.Eng., M.Sc. and Ph.D. degrees. There are approximately 160 undergraduate and 90 graduate students in the department, and excellent research facilities. Curriculum vitae, a list of referees and a statement of research interests should be addressed to: M.S. Sachdev, Head, Department of Electrical Engineering, University of Saskatchewan, Saskatoon, Canada, S7N 0W0. Applications must be received by November 5, 1992. The expected appointment date is January 1, 1993. The University of Saskatchewan is committed to the principles of employment equity. Canadian citizens and permanent residents will receive first consideration. Applications of non-Canadians will then be considered.

**Graduate Research Assistants.** The Center for Microelectronics Research (CMR) at the University of South Florida seeks graduate research assistants to pursue state-of-the-art research and related Ph.D. in the areas of microelectronic materials and defects, semiconductor processing and manufacturing, VLSI/ULSI/WSI architecture, circuit design, design automation, MCM design, and test. Successful applicants will be required to pursue Ph.D. in E.E. or Computer Science. Stipends available in \$12K to \$19K (half time) range for full year beginning Jan, May or Aug 1993. Tuition waivers available. Applicants must have excellent academic record and minimum of Bachelors degree in appropriate discipline. U.S. citizenship is a requirement for several positions. CMR research funding totals nearly \$14M over the past three years. USF is an AA/EO Employer. Resumes to Dr. Earl Claire, Director, CMR/USF College of Engineering, M/S ENG 118, 4202 E. Fowler Avenue, Tampa, FL 33620.

**Stanford University.** The Departments of Materials Science and Engineering and Electrical Engineering invite applications for a tenure track faculty position in the area of magnetic or optical information storage. A PhD and a strong interest in graduate and undergraduate teaching are required. The research activities of this individual should focus on the science and technology of information storage materials

and devices and might embrace magnetic thin films, magneto optic materials, optical storage materials, or other relevant materials and techniques. The successful candidate will be embedded in an environment with substantial activities in surface and interface science, magnetism, solid state physics, advanced optics and nanostructure science and technology and have an interest in the fundamentals and applications of both materials and devices. The opening is at the Assistant Professor level. The appointment will be made either in the MS&E Department or the EE Department or both. Please send a complete resume, a publication list, a statement of research and teaching interests, and the names of three references to Professor Robert L. White, Department of Materials Science and Engineering, MS 2205, Stanford University, Stanford, CA 94305 by December 1, 1992. Stanford University is an equal opportunity/affirmative action employer and encourages applications from women and minority candidates.

**Chairperson:** Electrical and Computer Engineering, New Jersey Institute of Technology. Candidate with proven research, teaching, program development and administrative record sought to provide leadership for the ECE department. NJIT is New Jersey's public technological university with nearly 7400 students pursuing baccalaureate through Ph.D degrees in Newark College of Engineering, School of Architecture, College of Science and Liberal Arts, and School of Industrial Management. The university expends approximately \$15 million annually in sponsored research, \$1 million of which is in the ECE department. Research areas include biomedical, communications, computer, and control systems, microwave and lightwave engineering, power systems and solid state circuits and devices. NJIT's campus is networked and computer intensive. Additional facilities include: interdisciplinary Center for Manufacturing Systems, state-of-the-art Microelectronics Center with class 10 cleanroom for integrated circuits fabrication, Center for Communication and Signal Processing, Center for Microwave and Lightwave Engineering, Optoelectronics Center, and modern offices and student laboratories. Enrollment is the highest of any ECE department in the New Jersey/New York metropolitan region; the department plays important roles in the profession and the community. Appropriate doctorate, solid credentials in research, publication and educational leadership required. Send resume to: Personnel Box ECEC. NJIT does not discriminate on the basis of sex, sexual orientation, race, color, handicap, religion, national or ethnic origin, veteran's status, lifestyle or age in employment. New Jersey Institute of Technology, University Heights, Newark, NJ 07102.

### Government/Industry Positions Open

**Development Staff Member:** Research and development of high capacity optical information storage systems involving new serial/parallel optical recording and readout techniques. Design and implement prototypes with full computer simulation-optical and systems predictive modeling. Address and enhance recording and readout device performance limits: noise sources, optical aberrations & interface electronics. Must have Ph.D. degree in Electrical Engineering, one yr. exp. in job offered or one yr. exp. as a pre or post doctoral research assistant/fellow. The one yr. exp. must include research in volume holographic information storage in real time photorefractive based optical data storage materials, including noise & crosstalk analysis, vector diffraction modeling, BPM diffraction algorithms, Fourier optics & digital signal processing as well as diode lasers, especially thermal & optical feedback effects, & noise generation issues. 40 Hrs/wk., \$57,200/yr. Qualified applicants send resume or application letter with ad by September 30, 1992 to: AZ DES Job Service, Attn: 732A, Re: 0017053, P.O. Box 6123, Phoenix, AZ 85005. Job Location: Tucson, AZ. Emp. pd. ad. Proof of authorization to work in U.S. required if hired.

**Technical Support Specialist.** For vendor of parallel supercomputer. Ports, modifies, tests & im-



## CLASSIFIED EMPLOYMENT OPPORTUNITIES

proves comp. fluid dynamics (CFD) programs. Helps develop num. methods & libraries to enhance performance. Familiarizes customers and prospects with company products and provides consultation and training in parallel programming. Collaborates with scientists and engineers and represents company at scientific conferences. Extensive travel & on-site support at customer locations. 40 hrs/wk. Salary \$60,000/yr. Based in Bellevue, WA. Req. Ph.D. in Aero. Engrg. or equiv.; 1 yr exp. in job offered or 3 yrs exp. as Sr. Res. Engr. and/or Univ. Res. Assoc. in CFD. Education must include univ. deg. in at least two different science or engrg. fields and include 2 yrs of grad. level course work in applied and comp. fluid dynamics. Exp. gained during education or prof. employment must include 1 yr in parallel/distributed processing; 3 yrs in CFD, num. methods and algorithms using Cray or comparable supercomputers on or off-site; 5 yrs in Unix, Fortran, C, shell programming, networking with TCP/IP, graphical user interfaces, and third-party applications software. Periods of exp. may overlap where appropriate. Must have proof of legal authority to work permanently in U.S. Send resume by Oct. 1, 1992 to Emp. Sec. Dept., E&T Div., Job# 319661-S, P.O. Box 9046, Olympia, WA 98507-9046.

**Software Engineer** for computer consulting company in central Ohio. Design, develop and integrate hardware and software elements to collect traffic data from telephone networks and process and deliver it to customers. Specific responsibilities include: design, develop and test front end data collection software in Unix C/C++ for real time environment using object oriented techniques; design, develop, and specify and port UNIX application and system level software, communication protocol and user interface, real time switching software; design, develop, enhance UNIX device drivers, firmware software, diagnostics for SCSI, serial asynchronous and synchronous communication for VME based RISC work station (Sun workstations) develop automation tools, test scripts, process and debugging tools for unit, regression and stress testing of the system; evaluate and finalize data communication equipment, multiplexors, SCSI, VME bus based peripherals, hardware installation, performance measurements, new feature design, development, coordination and management. Job requires Bachelor Degree in Electronics and Communication Engineering or Computer Science or Computer Engineering or Electronic Engineering and four years experience as Software Engineer or Systems Engineer. At least 1 year of experience must be in UNIX application and system level programming, porting and debugging in C, assembly language, Kernel enhancement, debugging and device driver development for RISC based work stations. At least one year of software development experience must be in data communication, interfacing and protocol development for serial, SCSI, multiplexors, communication networks and real time switching software. At least one year of experience must include front end data collection software development for VME based hardware, user interface, developing test scripts, debugging tools for hardware and software testing. Must have at least one undergraduate course in each of following: operating systems, communication systems, compiler design, data structures and network analysis and synthesis. 40 hrs/wk; 8 a.m.—5 p.m. Salary \$1,230.76/wk. Must have proof of legal authority to work permanently in United States. Send resume in duplicate (no calls) to J. Davies, JO #1260393, Ohio Bureau of Employment Services, P.O. Box 1618, Columbus, Ohio 43216.

**Technology Wanted.** Public Transportation Research Group seeks innovative technology to improve the efficiency of rail car and transit bus maintenance operations. Advanced technologies such as robotics and automated test procedures/equipment must be proven, cost-effective, and operational in other disciplines. Inquiries may be directed to: Transportation Development Corporation, 1201 New York Ave N.W., Washington, DC 20005.

**Position Available.** Engineer, Electromagnetic—

**Microwave Antennas—**Design, develop & test antennas and microwave components/systems for high gain tracking & spacecraft antennas. Duties: develop computer software for reflector geometry & electromag performance of symmetric, offset & multiple reflector ant.; analyze multimode waveguide feeds for gain, sidelobes, VSWR & cross polarization; prepare tech rpts & present briefings; req: PhD in Electromagnetics/Electrical Eng'g & 5 yrs exp in job; also exp w/Methods of Moments & sftwre dsgr; familiar w/Green's functions & circular coord's; Fourier-Bessel Transforms; exp in Unix environ & Fortran 77, C & Postscript. Salary: \$50,000/yr; Job site: El Segundo, CA. Send resume & copy of ad to: Job # DB 22032, P.O. Box 9560, Sacramento, CA 95823-0560, no later than Oct. 5, 1992.

**Switchgear Project Engineer:** superv const/assy highvoltage disconnect switches in Spokane, WA. Review prod design for compliance w/eng princ, co stds & cust specs. Eval & approve design chgs & specs from bluepr. Supv assy of ea unit. Dir proj mgr, assy supv & 10 laborers. Resp to meet/exceed test reqmts in accord w/Am Natl Std C37-30 and ANSI 37.24 by admin var tst prac. Req 10 yr exp in job offered. \$4,466/mo; 8-5pm; 40 hr/wk. Resume: Empl Sec Dept E&T Div, Job #327747, PO Box 9046, Olympia, WA 98507-9046 by 9/25/92.

**Graphic Systems Engineer (Western Region).** Responsible for all facets of scanner & system installations & service. Modify scanning systems from company & other equipment manufacturers to interface with DOS & Macintosh based computer systems for special applications & prepare technical documentation for use by technicians in installation, repair & upgrade of these systems. Perform electrical testing & analyses & supervise & train graphic systems technicians who install, repair & upgrade scanners. Schedule: 40 hrs/wk; 8:00am—5:00pm, Mon.—Fri. Requirements: B.S. (or equiv.) in Electrical Engineering. 2 yrs exp. in position offered or 2 yrs exp. as Field Service Engineer (Graphic Systems). Exp. may be contemporaneous with education and/or professional exp. & must include minimum 2 yrs.: (1) performing or supervising installation, commissioning, maintenance & repair of technologically advanced laser LED scanners & graphic arts systems; & (2) designing test apparatus & testing electronic equipment. Salary: \$42,500/yr. Jobsite: CA & Western Region. Interviews will be conducted by telephone or in person in CA or NJ at employer's expense. Send ad & your resume to: Job # LQ 6659, P.O. Box 9560, Sacramento, CA 95823-0560, no later than 10/1/92.

**Software Engineer:** Needed to design, code, and test C++ Objects and application environments. Requires Bachelors Degree in Computer Science and two years experience in 1) multiple GUI's i.e. Macintosh, MS-Windows, OS/2-PM, and X-Windows; 2) distributed objects and distributed computing; 3) developing reusable software components encapsulating code and data; 4) event driven programming; and 5) OOP using C/C++. Salary \$45,000/yr. Job Order No. 2827178. Contact Utah Job Service 5735 South Redwood Road, P.O. Box 11750, Salt Lake City, UT 84147-0750.

**Development Engineer.** Design digital/analog/microprocessor hdw that controls medical excimer laser coronary angioplasty syst. Coord. w/softw engineer to optimize hdw/softw tradeoff & minimize design effort. Help deter. product spec. Prep. Failure Mode and Effects Analysis. Construct prototypes. Help prod. staff w/pilot prod. phase. Supervise/audit printed circuit bd layout. Prep. material lists. Specify/prep. purchased-part documentation. REQS.: BSEE; 4 yrs exp. designing digital/analog/microprocessor circuits to operate in severe RFI environ. Incl. 1) documenting designs in accord w/FDA gd. manuf. practices; 2) design fail-safe syst.; 3) devel./perform tests to verify Failure Mode Effects Analysis, conduct software/hardware tests, document concept/design process, all in accord. w/FDA guidelines; 4) control RFI emissions. Working knowl. of 1) schematic capture/CAD; 2) programming Intel emulators/microcontrollers w/assembly lang; 3) DOS

in IBM-PC environ.; 4) excimer laser microcontroller interface reqs; 5) FDA/Center for Devices and Radiological Health reqs for microcontroller interface design/validation; 6) electronic interface reqs for connect. optical fiber syst. to excimer laser syst.; 7) manufacturing reqs of design to insure min. cost/max. reliability. 40 hr/wk. \$43,000/yr. Mail resume no later than Sept. 30, 1992 to: Job Order #CO455881, CO Dept of Labor & Employment, 600 Grant St., #900, Denver, CO 80203-3528, Attn: Employment Programs. Must provide proof of legal right to work in U.S.

**Senior Design Engineer:** Requires all course-work requirements toward a Ph.D. in Electrical Engineering, including research in multicarrier modulation, demodulation, finite length equalization, echo cancellation, and initialization; single carrier modulation, demodulation, equalization, and initialization; timing recovery, Fast Fourier Transform (FFT), adaptation to line conditions, Motorola digital signal processors (DSP) for implementation, DSP code development (including documentation from theory, to detail, to code), theory and practice of the interaction between hardware and analog elements, and digital code; and coursework in communication systems, data communications, VLSI implementation, error correcting codes, digital filters, and digital signal processing. To develop and code algorithms on digital signal processors for multicarrier and single carrier high speed modulation systems applied to copper-wire transmission media. \$72,000/yr. Job/Interview Site: Palo Alto, CA. Clip ad and send with resume no later than October 1, 1992 to IEEE Spectrum Magazine, P.O. Box 9-1, 345 E. 47th St., New York, NY 10017. Upon hire must show immediate ability to work in the United States.

**Electrical Engineers.** Leading consulting firm seeks licensed electrical engineers for Upstate New York Office. Qualified applicants must have P.E. and 8+ years of experience in commercial, industrial, and governmental building related projects. We offer excellent salary and benefits. Send detailed resume and salary history to: EMJ/McFarland-Johnson Engineers, Inc., P.O. Box 1980, Binghamton, New York 13902. EOE M/F/H/V.

## ISSA/NSA POLY

### USE YOUR "TICKETS" FOR FASTER CAREER GROWTH

Put our 27+ years experience placing technical professionals to work for you. All fees paid. Nationwide opportunities in Communications, Defense, Intelligence, Computer, Satellites and Analytical Sciences. If you earn over \$35,000, we have a better, more rewarding job for you ... right now. U.S. citizenship and ISSA/NSA POLY desirable. Call (301) 231-9000 or send your resume in confidence to: Dept. EA-13EB or FAX to: (301) 770-9015.

## WALLACH associates, inc.

Technical and Executive Search

Washington Science Center  
6101 Executive Boulevard  
Box 6016  
Rockville, Maryland 20849-6016



**SPECTRUM'S 1992 ENGINEERING/SCIENTIFIC SOFTWARE REPORT & GUIDE**

# SOFTWARE

**DESIGN, MATH, GRAPHICS AND MORE**

## WHAT'S UP IN ENGINEERING/SCIENTIFIC SOFTWARE?

### **ANNOUNCING SPECTRUM'S THIRD ANNUAL FOCUS REPORT & GUIDE TO ENGINEERING/SCIENTIFIC SOFTWARE-NOVEMBER 1992**

An update on what's new and what's vital for today's computer-aided design and engineering professional. Up-to-the minute tabulations that integrate key data and information on the most advanced engineering/scientific software packages for PCs and workstations.

Eminent authors and reviewers assess the advantages—and some disadvantages—among the following product categories:

- Logic synthesis for ASICs
- Mixed signal and analog design
- Electromagnetic simulation and design
- Multichip module routing and placement
- Data handling (acquisition, analysis, display and technical reporting)
- Math, graphics, visualization
- Development tools for embedded signal processing
- CAD frameworks

Plus two new, special features:

- Report on SPECTRUM's subscriber survey on software usage
- Report on users panel on framework utilization

Don't miss this exclusive, high profile marketing environment in our November issue. It's your opportunity to reach the core of the high-tech software market, over 320,000 SPECTRUM subscribers—100% engineering/scientific audience—the largest single concentration of electronics, communications, computer and high technology professionals in the world.

For more information, contact a SPECTRUM Regional Sales Office.

**Issue date November 1992****RESERVE SPACE NOW!****Ad closing date: October 1**

IEEE  
**SPECTRUM**  
M A G A Z I N E  
*The High-Tech Marketing Tool™*

345 East 47th Street, New York, NY 10017, U.S.A. 212-705-7760

Spectrum is published by The Institute of Electrical and Electronics Engineers, Inc.



## ADVERTISING RATES

1 Insertion—\$320

12 Insertions—\$3840

50% discount to IEEE members on three or more insertions.

If you are an IEEE member, please enclose your membership number with the order.

Copy cannot exceed 1-inch in depth.

No product may be offered for sale.

Advertising restricted to professional engineering and consulting services.

No orders can be accepted by telephone; order and copy must be sent together.

For any further information and closing dates please contact:

Advertising Production Dept., 212-705-7578.

## THE CONSULTING GROUP

Multi-Disciplined Engineers with P.E./Ph.D.

- Microwave, RF, Fiber-Optic Systems Design
- Oscillators, Amplifiers, Filters, Antennas, Synthesizer/PLL Design, Microprocessor, Communication cks. Industrial Power System.

R&D, Prototyping & Testing in our Lab facilities.

119-40 Metropolitan Ave., Kew Gardens, NY 11415

Ph. (718) 846-5400 Fax (718) 846-2440

## RAINES ELECTROMAGNETICS Consulting Since 1972

- Antennas and Arrays
- Scattering and Radar Cross Sections
- Radhaz & Environmental Impact
- Simulations of Fields & Phenomena

Jeremy K. Raines, Ph.D. (MIT), P.E.

13420 Cleveland Drive (301) 279-2972  
Potomac, Maryland 20850-3603

## CONTROL SYSTEM CONSULTING

- Servo design, high performance motion control, synthesis, system performance, simulation, specs, integration, testing
- Electrical, mechanical, hydraulic
- Defense, aerospace, industrial experience

30 East Gate Road A.R. Hazelton  
Danbury, Conn. 06811 (203) 743-7002

## IRA J. PITEL, Ph.D.

Consulting, Research and Development  
in Power Electronics and Magnetics  
Power Supplies, Inverters, Converters, Motor  
Drives, Lighting Controls, Industrial Controls,  
Transformers, and Special Magnetics.

**MAGNA-POWER ELECTRONICS, INC.**  
135 Route 10 Whippany, NJ 07981  
(201) 428-1197

## NOISE, TRANSIENTS AND INTERFERENCE

- FCC, VDE, EMC/EMI
- Susceptibility, ESD, RF, Transients, Lightning
- Testing and Retrofit for Product Enhancement
- Speed to Market via Design Review
- Noise-Immune Designs and Prototypes
- FCC Compliance Training and Retrofits

**TKC**  
(813) 544-2594

THE  
KEENAN  
CORPORATION

R. Kenneth Keenan, Ph.D.  
Vice President, Engr.  
8609 66th St. Pinellas Park, FL 33551

## LEONARD R. KAHN, P.E.

Consultant in Communications and Electronics  
Single Sideband and Frequency Shift Systems  
Diversity Reception - Stereophonic Systems  
Modulation Systems  
Registered Patent Agent

222 Westbury Ave.  
Carle Place, NY 11514  
516-222-2221

## International Compliance Corporation

### Design, Test, & Consulting

- FCC Certification/Verification
  - VDE, CISPR, VCCI (Japan)
  - "1992" European Compliance Testing
  - Product Safety: UL, CSA, IEC, VDE
  - Electrostatic Discharge (ESD)
  - MIL-STD 461/462, NARTE-Certified Engineers
- 1911 E. Jeter Rd (817) 491-3696  
Argyle, TX 76226-9401 Fax: (817) 491-3699

## IOCC

### INTEGRATED OPTICAL CIRCUIT CONSULTANTS

- Consulting, Contract R&D, and Prototyping
- Integrated, Fiber, and Guided-Wave Optics
- Applications Engineering
- Design, Fabrication and Evaluation
- Critical Analysis of Technology
- Troubleshooting
- Marketing

R.A. Becker, D. Sc. 10482 Chisholm Ave.  
President Cupertino, CA 95014  
(408) 446-9812

## T-TECH<sub>TM</sub>

1 DEAN ST./BOX 151  
HUDSON, MASS.

VOICE: (508) 562 5820 FAX: (508) 568 1219

RF (TO 2 GHz), ANALOG, VIDEO,  
FIBER OPTICS, PLL'S, SYNTHESIZERS,  
FILTERS, A/D CONVERSION

THE MOST TIMELY, COST EFFECTIVE,  
AND HIGHEST PERFORMANCE SOLUTION.

## Patent Attorney

Robert E. Malm, Ph.D. (M.I.T.)  
Attorney At Law

Post Office Box 522  
Pacific Palisades, CA 90272

Tel: (310) 459-8728  
Fax: (310) 573-1781

## SOFTWARE ENGINEERING

- Real-Time Systems Analysis and Design
- Software Engineering Training
- (Call for Offerings)
- CASE Training and Implementation
- Consulting Services
- Product Development

Carl A. Argila, Ph.D., Inc.  
SOFTWARE ENGINEERING CONSULTANT  
800-347-6903

## ELECTROMAGNETIC DESIGN ANALYSIS

Consultancy by world leaders in 3D electromagnetic  
computation using PE2D, OPERA, TOSCA and ELEKTRA

- electrical machines
- recording heads
- actuators
- transformers
- loudspeakers
- MRI scanners
- magnetic casting
- scientific apparatus
- NDT equipment
- accelerator magnets

VECTOR FIELDS INC. tel: (708) 851-1734  
1700 N. Farnsworth Ave, Aurora IL 60505 fax: (708) 851-2106

## CONSULTING & PROTOTYPES ELECTRIC MOTORS

BRUSHLESS MOTORS SWITCHED RELUCTANCE  
STEPPING MOTORS AC INDUCTION

MAGNA PHYSICS CORP. JAMES R. HENDERSHOT  
100 Homestead Ave. TEL: 513-393-9835  
P.O. Box 78 513-393-3810  
HILLSBORO, OH 45133 FAX: 513-393-9836

## INTEGRATED CIRCUIT DESIGN CONSULTING

Providing integrated circuit design and analysis  
consulting services with gate arrays, cell libraries, or  
full custom in digital, analog, cmos, bipolar or eel  
technologies and support services such as modeling,  
layout, verification, software, testing, debugging,  
documentation, and reverse engineering.

1556 Halford Ave., Ste 310, Santa Clara, CA 95051  
(408) 243-7422

## O.E.M. Electronic Products

- 13 Year Product Development History
- Custom, RF/Digital ASICs • RF Systems
- Computerized Instrumentation
- Spread Spectrum Communications
- Hitachi Authorized Design Center

## LocUS, Incorporated

1842 Hoffman St., Madison, WI 53704  
608/244-0500 FAX 608/244-0528

## Princeton Electro-Technology, Inc.

**PERMANENT MAGNETS  
MATERIALS, DEVICES & APPLICATIONS**  
CONSULTING, DESIGN, PROTOTYPES, MARKETING

Peter Campbell, Ph.D., President  
2449 Patricia Avenue Tel: (310) 287-0375  
Los Angeles, CA 90064 Fax: (310) 287-0378

## INDUSTRIAL CONTROLS

### CUSTOM MICROPROCESSOR CONTROLS

- Machine Tool Controls
- Process Controls
- Medical Electronics
- Neural Networks
- Consumer Product Design
- Instrument Design
- Prototype Production
- Control Software

**W** Wintriss Engineering Corporation  
4715 Viewridge, San Diego 92123  
(800) 733-8089 Vic Wintriss, MSEE

## OSI MAP GOSIP

- Education & Newsletters
- Market & Product Planning
- CIM Plans, Design & Integration

### SHIP STAR Associates

Bob Crowder, President  
36 Woodhill Drive, #19 Newark, DE 19711  
Tel: 302-738-7782 FAX: 302-738-0855

**802 FIELDBUS MMS**

## Patent Attorney

**Electronics-Signal Processing**

I. Ionescu, MSEE, PE, JD

Patent Office Reg # 33185

203-381-9400 (voice)

203-381-9401 (fax)

Box 697, Stratford, CT 06497



## European Patent Attorney

ing. Umberto MONTI  
Via Washington 48  
I - 20146 Milano Italy  
Filing and prosecution of  
PCT & European Applications

## Communications and Control Systems

- Real-Time Analysis, Design and Implementation
- Process Control Systems and Automated Testers
- Store and Forward, Circuit Switching, and DSP
- Large Variety of Processors, Languages, OS
- Defense and Commercial Industrial Sectors
- DoD 2167A and NSAM 81-3 Development

### CommSys Software Engineering

Tel: (609) 234-8088 Fax: (609) 234-0323

DAVID NEWMAN & ASSOCIATES PC  
PATENT AND TRIAL LAWYERS

PhD, Elect Engr, JD

## PATENT, COPYRIGHT AND COMPLEX TECHNOLOGY RELATED CASES

\*\*\* NATIONWIDE \*\*\*

PROSECUTION • LITIGATION

TEL 301 934-6100 FAX 301 934-5782

## SSPI

WILLIAM A. GARDNER, President  
*Specialists in Exploitation of Cyclostationarity  
in Signal Processing for Detection,  
Classification, Location, and Extraction*  
6950 Yount Street, Yountville, CA 94599  
(707) 944-0648 Fax: 944-0144

## SPREAD SPECTRUM

- Communications Systems Engineering, Inc.  
• Specializing in Spread Spectrum systems.  
• FCC Part 15 apps, direct sequence or hopping.  
• Low cost, alignment free implementations, DSP.  
• Complete RF and digital laboratory facilities.

1004 Amherst Avenue Phone: (310) 820-3825  
L.A., CA 90049 Fax: (310) 820-6761

## PII FORCEFIELD INDUSTRIES, INC.

MAGNETIC, FERROELECTRIC & FERROIC  
MATERIALS, DEVICES & APPLICATIONS  
CONSULTING, RESEARCH & PROTOTYPES

Frederick Rothwarf, Ph.D., President  
John Popplewell, Ph.D. Dilip Das-Gupta, Ph.D.  
11722 Indian Ridge Road, Reston, VA 22091, USA  
(703) 758-0247 (703) 620-1784 (FAX)  
17 Trefonwys, Bangor, Gwynedd LL572HU, U.K.  
044-248-353607

## EMBEDDED CONTROLLERS

Specializing in the Intel/Signetics 8051 Family

- System design
- Hardware design
- Hardware de-bugging
- "C" and assembly language programming
- Code size reduction
- Protocol development/implementation

Amalgamated Engineering eMail (CIS) 71064,104  
Burlingame, California voice (415) 375-1988  
fax (415) 579-4688

## CONTINGENCY PATENT LICENSING NO FEES ANYTIME

Three decades experience: Law, Technology, Negotiation  
Five Earned Degrees

### DAVID FINK, ESQ.

441 Summer Street  
Stamford, CT 06901  
Tel: (203) 325-3344 Fax: (203) 325-4443

## MATERIALS RESEARCH, INC.

PO BOX 566 • CENTERVILLE, UT 84014 • (801) 298-4000

### 18 Years of Research Providing:

- PRODUCT & PROCESS INNOVATION
- SURFACE ANALYSIS FROM -190°C to 1200°C
- NEW MATERIALS FORMULATION
- FAILURE ANALYSIS
- MANUFACTURING SOLUTIONS

## BOARD OF DIRECTORS

Merrill W. Buckley Jr., *President*  
Martha Sloan, *President-Elect*  
Theodore W. Hissey Jr., *Treasurer*  
Karsten E. Drangeid, *Secretary*

Eric E. Sumner, *Past President*  
Eric Herz, *Executive Director*

## Vice Presidents

Edward A. Parrish, *Educational Activities*  
Arvid G. Larson, *Professional Activities*  
James T. Cain, *Publication Activities*  
Luis T. Gandia, *Regional Activities*  
Marco W. Migliaro, *Standards Activities*  
Fernando Aldana, *Technical Activities*

## Division Directors

Kenneth R. Laker (I)	V. Thomas Rhyne (VI)
Lloyd A. Morley (II)	Robert A. Dent (VII)
Frederick T. Andrews (III)	Helen M. Wood (VIII)
Martin V. Schneider (IV)	Jan Brown (IX)
Bill D. Carroll (V)	H. Vincent Poor (X)

## Region Directors

Joel B. Snyder (1)	Jerry C. Aukland (6)
Charles K. Alexander Jr. (2)	Vijay K. Bhargava (7)
David A. Conner (3)	Kurt R. Richter (8)
Howard L. Wolfman (4)	Eduardo Arriola (9)
James V. Leonard (5)	Souguil J.M. Ann (10)

Donald G. Fink, *Director Emeritus*

## HEADQUARTERS STAFF

John H. Powers, *General Manager*  
Eric Herz, *Executive Director*

## Associate General Managers

Thomas W. Bartlett, *Finance and  
Administration*  
William D. Crawley, *Programs*

## Staff Directors

Donald Christiansen, *IEEE Spectrum*  
Irving Engelson, *Technical Activities*  
Leo C. Fanning, *Professional Activities*  
William Habingreither, *Customer Service*  
Phyllis Hall, *Publishing Services*  
Melvin I. Olken, *Field Services*  
Andrew G. Salem, *Standards*

## Staff Secretaries

Awards: Maureen Quinn  
Board of Directors: Mercy Kowalczyk  
Educational Activities: Peter A. Lewis  
Regional Activities: Melvin I. Olken  
Publishing: Phyllis Hall  
Standards Activities: Andrew G. Salem  
Technical Activities: Irving Engelson  
United States Activities: Leo Fanning  
For more information on Staff Secretaries  
to IEEE Committees, please communicate  
with the IEEE General Manager.

## PUBLICATIONS BOARD

James T. Cain, *Chairman*  
Lloyd A. Morley, *Vice Chairman*  
Phyllis Hall, *Staff Secretary\**

Charles K. Alexander, Donald  
Christiansen\*, Sajjad H. Durrani, Irving  
Engelson, Randall L. Geiger, Leo Grigsby,  
Abraham H. Haddad, Ronald G.  
Hoelzeman, Tatsuo Itoh, Ted Lewis, Donald  
R. Mack, William Perkins, G.P. Rodrigue,  
Daniel Rosich, Allan C. Schell, Leonard  
Shaw, Sol Triebwasser, Robert B. Voller,  
Stephen B. Weinstein, Ronald D. Williams  
\*Ex officio

## Publishing Services

Phyllis Hall, *Staff Director*  
Jim Ashling, *Electronic Publishing Products  
Director*  
Patricia Walker, *Magazines Director*  
Ann H. Burgmeyer, Gail S. Ferenc, *Managers,  
Transactions*  
W. Reed Crone, *Managing Editor, Proc. IEEE*  
William Hagen, *Administrator, Copyrights  
and Trademarks*  
Dudley Kay, *Managing Editor, IEEE Press*  
Lewis Moore, *Business Manager*  
Adam D. Philippidis, *Manager,  
Indexing and Abstracting*  
Eileen Wilson, *Manager, Special Publications  
and Publishing*

Editorial Offices: New York City 212-705-7555 San Francisco 415-282-3608 Washington, DC 202-544-3790

EDITORIAL CORRESPONDENCE. *IEEE Spectrum*, 345  
East 47th St., New York, N.Y. 10017, Attn: Editorial Dept.  
Responsibility for the substance of articles published rests upon  
the authors, not the IEEE or its members. Letters to the editor  
may be excerpted for publication.

REPRINT PERMISSION. Libraries: Articles may be photo-  
copied for private use of patrons. A per-copy fee (indicated  
after the code appearing at the bottom of the first page of  
each article) must be paid to the Copyright Clearance Center,  
29 Congress St., Salem, Mass 01970. Instructions: Isolated ar-  
ticles may be photocopied for noncommercial classroom use

without fee. Other copying or republication: Contact: Editor,  
*IEEE Spectrum*.

ADVERTISING CORRESPONDENCE. *IEEE Spectrum*, 345  
East 47th St., New York, N.Y. 10017, Attn: Advertising Dept.,  
212-705-7760. Also see Advertising Sales Offices listing on last  
page. The publisher reserves the right to reject any advertising.  
COPYRIGHTS AND TRADEMARKS. *IEEE Spectrum* is a re-  
gistered trademark owned by The Institute of Electrical and  
Electronics Engineers, Inc. EE's tools & toys, Faults & failures,  
Innovations, Legal aspects, Managing technology, Newslog,  
Program notes, Reflections, Speakout, Spectral lines, Spinoffs,

Technically speaking, The engineer at large, and Whatever hap-  
pened to? are trademarks of the IEEE.

GENERAL INQUIRIES: 1-800-678-IEEE; Headquarters:  
345 East 47th St., New York, N.Y. 10017, 212-705-7900.  
Tel. extensions at headquarters: General Manager, 7910;  
Public Relations, 7369; Publishing Services, 7560. At other  
locations: Regional Activities, Section Services, Educational  
Activities, 908-981-0060, or write to IEEE Service Center, Box  
1331, Piscataway, N.J. 08855; Membership Services, 908-  
981-0060; Professional Services, 202-785-0017; Technical  
Activities, 908-562-3900.



# Scanning The Institute

## Belt-tightening for all

At its meeting in early August, the IEEE Board of Directors approved a general operating budget for 1993 of US \$55.9 million that allows for a deficit of \$189 000. The final budget called for some severe belt-tightening after the initial budget drawn up in May called for a deficit of \$1.158 million.

"We faced a difficult task in coming as we did to a degree of balance," said Treasurer Ted Hissey. "We did it by challenging the staff and volunteers to come up with a sharp reduction in expenses to go along with a dues increase."

To close the gap, the amount originally budgeted for staff expenses next year was reduced by \$550 000, and the figure for travel by both staff and volunteers was reduced by \$450 000. There were also a number of increases in member dues and Regional assessments. Dues increased:

- \$5 to a total of \$78 for all members.
- \$5 to a total of \$28 (of which \$8 is for a subscription to *Potentials* magazine) for students in Regions 1-7 and to \$2 for Regions 8-10.

- \$1 to a total of \$33 for affiliates.

Some Regional assessments were also increased:

- \$2 to a total of \$24 for Regions 1-6.
- \$3 to a total of \$15 for Region 7.
- \$1 to a total of \$3 for Region 10.

## Strategic retreat

A cross section of IEEE leadership took part in late July in a first-of-its-kind strategic planning retreat. Their task: to consider the process of making strategic plans for the Institute out as far as 10 years.

Strategic planning has been handled by the Strategic Planning Committee, which mainly produced reviews of the plans made by individual Boards. It made recommendations "in reports that never seemed to get out of the file cabinet," pointed out IEEE President Merrill W. Buckley Jr.

Purpose of the retreat was to get strategic planning "into the mainstream," said Buckley. Accordingly, 47 people attended the retreat, including members of the Strategic Planning Committee, Executive Committee, major Boards, and Transnational Committee, as well as both 1993 presidential candidates, and key staff executives.

Divided in teams, participants dealt with a range of issues, including the global market and information environments in which the Institute will be operating, and issues related to education, technology, and human relations. Also considered were finances, membership, and products and services.

The teams were to make recommendations in September to the Strategic Planning Committee, which will organize them into a cohesive report. Follow-up reports are due to the Executive Committee in October and the Board of Directors in December.

## Metric policy

The IEEE Standards Board has formed an *ad hoc* Committee on Metric Policy, chaired by Bruce B. Barrow (F), to develop an IEEE-wide policy on metric usage. The senior boards and committees of the IEEE are invited to appoint individuals to the new committee. And any member with a point of view can make it known by writing to Cathy Goldsmith, IEEE, Box 1331, Piscataway, N.J. 08855-1331.

## Coming in Spectrum

**ENTER E-MAIL.** Electronic mail is gaining ground everywhere, not only in the workplace but also for recreational purposes. This special report includes an overview of the field, discussions of applications, and the phenomenon as a force for social change.

**SUPERFAST DYNAMIC RAMs.** RAMs have lagged far behind microprocessors in operating speeds, complicating system design. Now a new breed of superfast DRAMs will make a single memory system possible for many new computer applications. The emerging industry standards, chips, and data-bus architectures are discussed in this special report.

**SYSTEMS ARCHITECTING.** This difficult art copes with ill-structured, unbounded problems. It involves meshing systems into an end result in compliance with certification needs, public policy, and post-sale accountability.

**FUZZY DESIGN.** A general methodology for designing fuzzy logic systems advises on how to define a model, segment it into fuzzy and other components, isolate critical variables—and pin down the "pathological state."

**PAN-EUROPEAN POWER.** With the opening up of Eastern Europe, most of the continent is working to integrate its two great electric-power grids into a single system.

**MINORITY ROLE MODEL.** Since World War II, Howard Jones has earned 31 patents in radar and antenna technology. Now retired, the black engineer is in demand both as a consultant and as a role model for budding minority engineers.

## ADVERTISING SALES OFFICES

345 East 47th Street, New York, N.Y. 10017

William R. Saunders	Advertising Director	212-705-7767
Michael Triunfo	Northeastern Manager	212-705-7312
Dawn Becker	New York Manager	212-705-7311
Hendrik Prins	Research Manager	212-705-7067
Theresa Fitzpatrick	Advertising Production Manager	212-705-7579

Boston, Mass.—Doug Long,  
46 Main St., P.O. Box 2780, Orleans, 02653 508-255-4014

Chicago, Ill.—Ralph Bergen,  
One Northfield Plaza, Suite 300, Northfield, 60093 708-446-1444

San Francisco, Calif.—Denis F. Duffy,  
3755 Balboa Street, Suite 201, 94121 415-386-5202

Los Angeles, Calif.—Lynne Andrew Evans,  
5757 West Century Boulevard, #700, 90045 310-849-3800

Atlanta, Ga.—H. Michael Coleman, C. William Bentz III  
4651 Roswell Road N.E., 30342 404-256-3800

Dallas, Texas—Matt Kincaid, Eric Kincaid,  
9794 Forest Lane, Suite 634, 75243 214-553-9896

Tokyo, Japan—German Tajiri  
IMI Corp., Saito Lend Bldg., 13-5, Ginza 3-chome, Chuo-ku,  
Tokyo 104. (03) 3546-2231. TELEX: J23449 (IMI TOKYO)

Munich, Germany—Ric Bessford, Leopoldstr. 52,  
8000 Munich 40, Germany, 49-89-39-00-55

Publisher: Donald Christiansen  
Administrative Assistant: Nancy T. Hantman  
Associate Publisher: William R. Saunders (Advertising)  
Administrative Assistant: Carmen Cruz  
Vice President, Publications: James T. Cain

## ADVERTISERS INDEX

Circle numbers on the Reader Service Card, opposite page 80, that correspond to the advertisers listed below. \*Advertiser in North American edition.

RS#	Advertiser	Page #
50	*Athena Group	86A
30	*Aware	14H
54	*Beige Bag Software	86A
37	*Bimillennium	14A
	CACI	1
15	Computer Associates	23
24	CSPI	14
34	David Sarnoff Research Center	49
51	*Detronics	86A
1	DSP Development	cover 4
35	*Electroswitch	14J
39	*Electroswitch	14D
32	Global Intuition Network	79
38	*Granville Sewell	14H
	IBM	15
	IBM	24
	IEEE Spectrum	91
23	IEICE	52
43	*Integral Signal Processing	14D
18	Integrated Systems	2
5	Intusoft	83
6, 7	Jandel Scientific	7
27	Jandel Scientific	78
	*Japan New Materials Report	86A
36	*Keithley Metrabyte	14G
16	MacNeal-Schwendler	12
26	Macsyma	85
8	Mathsoft	13
4	Mathworks	50-51
40, 41,		
42	Microsim	11
17	Micro System Technologies 92	87
19	*Microway	14 I
10, 11,		
12, 13	Mita Press	16
14	Momentum Data System	cover 3
	*Motorola	14B-C
25	MRJ	80
22	National Instruments	18
53	*Polysoft	86A
31	Precision Visuals	cover 2
	*Seabury & Smith Insurance	21
9	Scientific Endeavors	86
21	Structural Research & Analysis	8
33	Superconductor Technologies	77
3	*Technical Career Institute	86F
20	Wolfram Research	17
52	*Z-World	86A



# MOMENTUM UNLOCKS INTEGRATED DSP SOLUTIONS

## GET YOUR KEY FOR UNLOCKING POWERFUL EASY-TO-USE DSP SYSTEMS

### Presenting the premier products from Momentum Data Systems

*QEDesign™* gives you the key to designing the most basic to the most complex digital filters that you'll ever need.

*DSPworks™* helps you unlock DSP solutions with its real-time data acquisition and signal processing.

*Momentum Code Generators* lock up your winning edge. Designed to create highly optimized assembly language programs with QEDesign.

Momentum products save you time and money with their menu-driven and easy-to-use solutions. No long learning curve or complex commands to memorize.

Momentum is dedicated to serving its customers by unlocking integrated DSP solutions. Whether its software only, or a combination of both hardware and software, Momentum has the key.

Momentum has software and hardware products available on SUN SPARCstations, Macintosh, PC's and Windows 3.0, and offers complete sales, service, training and consulting services.

Use the Momentum Key to unlock the best DSP solutions available from one company. Call today for a complete demo kit.

**714-557-6884**

**Now Available on Sun 4**



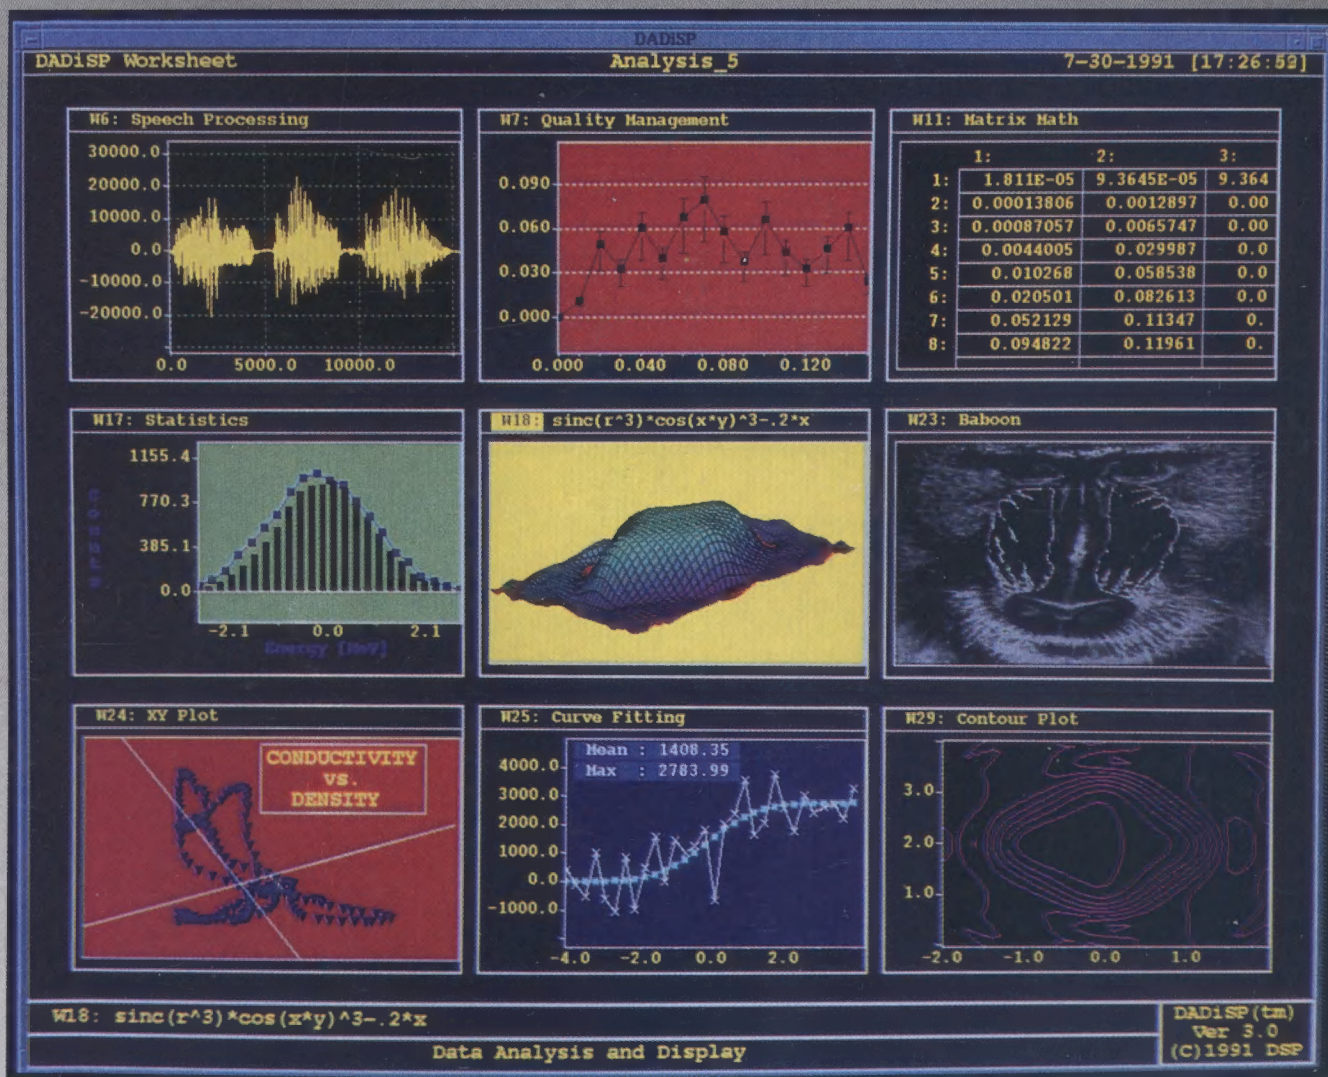
  
**MOMENTUM  
DATA SYSTEMS<sup>INC.</sup>**

1520 Nutmeg Place, Suite 108, Costa Mesa, CA 92626, U.S.A.  
Ph (714) 557-6884 Fax (714) 557-6969

**Australia:** SPA (Pty) Ltd, Tel: (3) 729-0995, Fax (3) 720-4298; **Denmark:** Assentoft Electronics, Tel 86 16 29 26, Fax 86 16 20 12; **France:** Excem, Tel (1) 47 52 13 44, Fax (1) 47 77 03 43;  
**Germany:** Electronic Tools, Tel (2102) 88010, Fax (2102) 880123; **India:** Signion Systems, Tel (842) 38085; **Italy:** CAD Italia Srl, Tel (2) 617 2521, Fax (2) 612 0003  
**Switzerland:** IDK, Tel (56) 27 27 77, Fax (56) 27 27 59; **United Kingdom:** Loughborough Sound Images Ltd., Tel (0509) 231843, Fax (0509) 262433; **Israel:** IES SAAR Ltd., Tel (972) 3 7526333,  
Fax (972) 3 7510927; **Japan:** LinX, Tel (81) 354893871, Fax (81) 354893872 Product names are trademarked to their respective companies.



# The Big Picture in Visual Data Analysis



## DADiSP 3.0

*The Standard Worldwide —  
adopted by thousands of scientists and  
engineers for data analysis and graphics.*

### Just Point-and-Click...

DADiSP works the way you do — from input to output. DADiSP's point-and-click menus put *you* in control: from Data Collection, Reduction and Management, through Visual Data Analysis and Hypothesis Testing, on to final output with DADiSP's *outstanding* Presentation Quality Graphics.

DADiSP is currently in use in engineering, laboratory data collection, matrix processing, manufacturing, science, signal processing, chemical and mechanical applications in automotive, aerospace, defense, medical, and other industries.



**CALL  
1-800-777-5151**

for your free DADiSP Trial Kit. DADiSP is available for SUN, HP, IBM, NeXT, DEC, Concurrent, and Silicon Graphics workstations, and of course, IBM PC compatibles.



One Kendall Square, Cambridge, MA 02139  
617-577-1133, FAX: 617-577-8211

Australia-Interworld Electronics, (03) 563-7066; Belgium-Electronique-Mesures, (02) 384-7309; Denmark-Engberg, (042) 25-1777; Finland-Mespek OY, (9) 0-351-1800; France-SM2i, (01) 34-89-78-78; Germany-Digi-Logger, (07720) 38022; Informatik Systeme, (0631) 35119-0; Ingenieurburo Kohler, (069) 9765-000; PDV-Systeme, (05321) 80761; Hungary-Selectrade Computer, (01) 163-2905; India-Dynalog Microsystems, (022) 517-0514; Israel-Racom, (03) 491922; Italy-BPS Computers, (02) 612-90221; BRM Italiana, (011) 771-00-10; Japan-Astrodesign, (044) 751-1011; Korea-Echo Microsystems, (02) 704-5378; Netherlands-Computer Engineering Roosendaal, (01650) 57-417; New Zealand-GTS Engineering, (09) 3092-464; Spain-Novatronic, (04) 452-0811; Sweden-SYSTEC, (013) 11-01-40; Switzerland-Urech & Harr, (061) 61-13-25; Taiwan-Advantech, (02) 351-2117; U.K.-Adept Scientific, (0462) 480055